TransSino: Prior Sinogram Pattern-Based Transformer for Limited-Angle CT Image Segmentation

Jae Hyun Yoon, Yeong Jong Lee, and Seok Bong Yoo*

Department of Artificial Intelligence Convergence, Chonnam National University, Gwangju, Korea
sbyoo@jnu.ac.kr

Abstract. Medical diagnosis using limited-angle computed tomography (LACT) is a beneficial approach for patients due to advantages such as faster scanning times and lower radiation doses. However, images reconstructed from LACT contain limited information, leading to significant artifacts and making an accurate diagnosis more challenging. Although various methods have been proposed to reconstruct LACT images into full-angle computed tomography (CT) images, they primarily focus on improving image quality and operate independently of lesion segmentation models, neglecting critical lesion-related information. In this paper, we propose TransSino, a transformer-based medical image segmentation model that operates in the sinogram domain of LACT. TransSino learns an auxiliary task to reconstruct the unmeasured regions in the sinogram domain for robust segmentation performance. Specifically, it analyzes the sequential nature of the sinogram using the transformer from language models and reconstructs features for the unmeasured regions by using prior sinogram patterns. Moreover, we introduce a contrastive abnormal feature loss to enhance the contrast between abnormal and normal feature regions. Experimental results confirm that TransSino outperforms existing medical segmentation methods on LACT images. The code is available at https://github.com/jhyoon964/TransSino.

Keywords: Limited-angle computed tomography · Medical image segmentation · Sinogram domain · Language models.

1 Introduction

Computed tomography (CT) is a widely used imaging technique for medical, industrial, and security applications due to its ability to provide images of internal structures. As illustrated in Fig 1(a), a CT scanner operates by emitting X-rays from a source toward the target object and measuring the attenuated signals with a detector to generate projections. Rotating 360 degrees, it sequentially

* Corresponding author

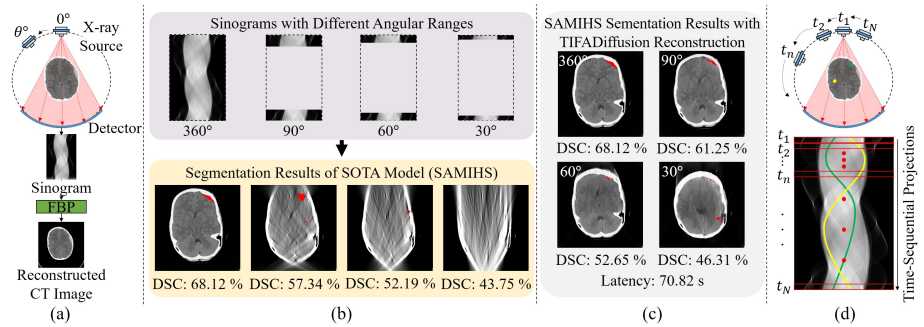

Fig. 1. (a) Acquisition of sinogram and CT image. (b) Visualization of the sinogram, reconstructed CT images, and the corresponding performance variations of the existing SOTA medical segmentation model according to the angular range. (c) Performance variations of the segmentation model with existing SOTA LACT reconstruction. (d) The CT projections acquired in time sequence and sinusoidal characteristics. The yellow and green sinusoidal patterns correspond to the sinogram from point in the object.

acquires these projections, which are aggregated into a sinogram. The sinogram is then processed using algorithms like filtered back projection (FBP) to reconstruct the CT image, enabling physicians to analyze internal anatomy for diagnosis and treatment.

Despite its advantages, CT scanning raises concerns due to the radiation exposure, which poses risks to patients requiring frequent scanning and those vulnerable to radiation, such as pediatric or cancer patients. Limited-angle CT (LACT) offers an alternative by reducing the angular range of projections, significantly decreasing the radiation dose. Additionally, LACT enables faster scanning by restricting the physical movement of the scanner and supports applications with physical constraints that limit scanning angles [1,5,7]. However, as depicted in Fig. 1(b), images reconstructed from LACT suffer from artifacts caused by the limited projection data, resulting in degraded image quality and obscured anatomical details. As the angular range becomes more restricted, these artifacts become increasingly severe, making an accurate diagnosis more challenging. The bottom row of Fig. 1(b) presents the segmentation performance, measured by the Dice similarity coefficient (DSC), of the existing state-of-the-art (SOTA) model for intracranial hemorrhage (ICH), SAMIHS [25], on the CTICH dataset [9] as the angular range is reduced from 360° to 30°. The results indicate that DSC declines due to artifact impact, presenting the limitations of existing segmentation models [8,10,13,19,22,27] under LACT conditions.

Many efforts to reconstruct LACT images into full-angle CT images have been explored to improve image quality. Recently, advanced deep learning-based models [11,12,18,23,26,28] have been proposed, demonstrating effective artifact removal and visually superior results. However, these existing LACT reconstruction models primarily focus on improving visual quality, often overlooking critical diagnostic information. As shown in Fig. 1(c), segmentation results from

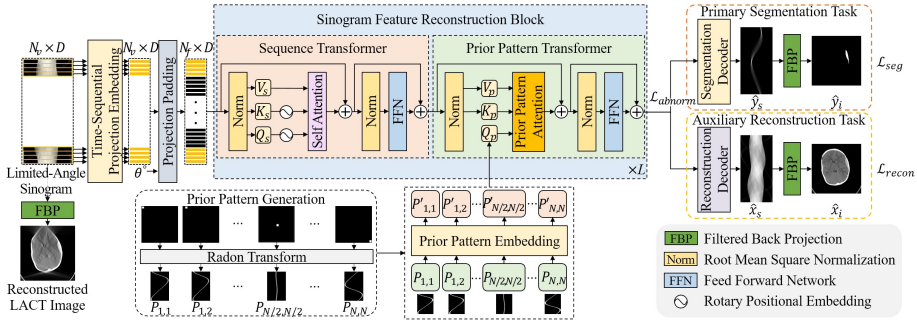

Fig. 2. Overall architecture of the proposed TransSino.

SAMIHS on CT images reconstructed using a SOTA LACT model, TIFADiffusion [26], reveal this limitation. Although the visual quality of the images is improved, DSC remains suboptimal, as key lesion information is inadequately recovered. Moreover, as shown in Fig. 1(d), the sinogram is acquired sequentially in terms of angle and time, and exhibits combination of continuous sinusoidal patterns. However, existing models do not account for these sequential and sinusoidal characteristics. Further, the latency of sequentially combined models is measured at 70.82 seconds, which presents a challenge for achieving rapid diagnosis in real-world clinical settings. These problems reflect the limitations of existing reconstruction methods that operate independently of segmentation models and fail to use the characteristics of the sinogram, which contains intact information before reconstruction via algorithm such as FBP.

To address these problems, we propose TransSino, a transformer-based medical image segmentation model using sinogram, which is unified with an auxiliary reconstruction approach and designed to be robust under LACT conditions. Our main contributions are as follows: **(i)** We propose a transformer-based segmentation model unified with an auxiliary reconstruction approach. Inspired by language models, this model handles the time-dependent characteristics in the sinogram for LACT. **(ii)** We propose a prior pattern transformer module to reconstruct unmeasured projections using the prior sinusoidal pattern cues of the sinogram with disentangled representation. **(iii)** We propose a contrastive abnormal feature loss that contrasts the characteristics of abnormal lesion regions with normal non-lesion regions in the feature space. This approach is achieved using paired CT data generated via an inpainting-based data augmentation technique.

2 Methods

As depicted in Fig. 2, we propose TransSino, a robust segmentation model designed for LACT conditions. TransSino embeds time-sequential projections and applies projection padding for unmeasured regions. Afterward, a sequence transformer captures temporal relationships, while a prior pattern transformer re-

stores unmeasured projections using prior sinusoidal patterns. The refined sinogram features are fed into segmentation and reconstruction decoders, guided by the total loss $\mathcal{L}_{total}$, which integrates segmentation, reconstruction, and abnormal feature losses ($\mathcal{L}_{seg}$, $\mathcal{L}_{recon}$, and $\mathcal{L}_{abnorm}$) for enhanced performance.

2.1 Sinogram Feature Reconstruction

As mentioned in Fig. 1(d), the sinogram has time-sequential and sinusoidal characteristics, which inherently contains the temporal and angular relationships, indicating the applicability of language models. Although existing models use the sinogram as additional information, they often overlook its time-sequential nature and do not properly handle its complex entangled sinusoidal patterns, leading to suboptimal reconstruction outcomes. To address this limitation, we introduce a sinogram feature reconstruction block to analyze the sinogram's inherent features and enhance restoration performance.

As shown in Fig. 2, we use sinogram data to handle the time-dependent characteristics of the CT projections. Although we do not directly utilize natural language, we adopt the transformer architecture inspired by language models such as Llama [6], treating each of the projections in CT as similar to word tokens. Each of the N_v projections passes through a time-sequential projection embedding, consisting of a linear layer, to obtain embedded projections with dimensions $N_v \times D$. The embedded projections along with angle information θ° are then fed into a projection padding module, which creates space for the unmeasured projections and aligns the dimensions with a full-angle sinogram. The resulting projections are used as input to the sinogram feature reconstruction block, consisting of sequence transformer and prior pattern transformer. The sequence transformer and prior pattern transformer are alternated over L cycles, progressively refining the features to extract the final sinogram feature f. Then, for auxiliary learning, f is fed into two decoders: the primary decoder for segmentation and the auxiliary decoder for reconstruction.

Sequence Transformer. Each embedded projection undergoes root mean square normalization (RMSNorm) and is passed through a linear layer to generate the query Q_s, key K_s, and value V_s. The Q_s and K_s are further processed through a rotary positional embedding and then enter the self-attention mechanism inspired by Llama [6]. The resulting feature is added to the previous output through a residual connection, followed by RMSNorm and a feed forward network (FFN). This process enables the model to capture the relationships among projections and infer projection information for the unmeasured angular regions.

Prior Pattern Transformer. As shown in the bottom left of Fig. 2, the pregenerated sinograms exhibit N^2 sinusoidal patterns as a result of the Radon transform [21] applied to $\frac{H}{N} \times \frac{W}{N}$ binary blocks shifted across the $H \times W$ image. Each point in the image is mapped to a distinct pattern in the sinogram. Based on this observation, we introduce the prior pattern transformer to restore the patterns in the unmeasured regions of the sinogram by using the prior knowledge.

While the sequence transformer focuses on analyzing the relationships among projections, the prior pattern transformer captures the relationships between the

distinct sinusoidal patterns within the sinogram. In the sinogram obtained from CT scanning, these sinusoidal patterns are intricately intertwined, making them challenging to interpret. Inspired by disentangled representation learning [24], we use prior patterns, which facilitate disentangling complex patterns in the sinogram and guide the generation of unmeasured regions. For this, the prior pattern transformer separates the sinogram feature obtained from the sequence transformer into N^2 channels. It then analyzes the correlation between the separated N^2 channels and the pre-generated N^2 prior patterns through an prior pattern attention mechanism. This mechanism captures the relationships between the distinct sinusoidal signals within the sinogram. The resulting feature is added via a residual connection and further processed via RMSNorm and FFN.

2.2 Auxiliary Learning of Segmentation and Reconstruction

Existing LACT reconstruction models focus on producing visually superior results over the global image but often fail to account for lesion-related local regions. This limitation arises because these models operate independently of segmentation models, forcing the segmentation task to rely on potentially unstable outputs. To address this issue, we introduce an auxiliary learning approach unifying the segmentation and reconstruction tasks.

The reconstruction decoder consists of an RMSNorm, a linear layer, and two convolutional layers, outputting the restored sinogram $\hat{x}_s$. The $\hat{x}_s$ is then passed through FBP to reconstruct CT image $\hat{x}_i$. For training, we define the reconstruction loss as follows:

$$\mathcal{L}_{recon} = \|x_s - \hat{x}_s\|_2 + \|G_x(x_s) - G_x(\hat{x}_s)\|_1 + \|G_y(x_s) - G_y(\hat{x}_s)\|_1 + \|x_i - \hat{x}_i\|_2, \quad (1)$$

where x_s and x_i are the ground truth (GT) of sinogram $\hat{x}_s$ and CT image $\hat{x}_i$. G_x and G_y denote the gradient operators for horizontal and vertical directions.

The segmentation decoder and loss $\mathcal{L}_{seg}$, consisting of Dice and cross-entropy losses, follow TransUNet [3]. The decoder output is a sinogram $\hat{y}_s$, which is passed through FBP to produce the segmentation map $\hat{y}_i$. Unlike existing models operating in the image domain, TransSino directly analyzes the sinogram, enabling the application of language models that capture its sequential nature. Further, it preserves critical signals that may be lost during FBP reconstruction, allowing for more precise segmentation. Then, using $\hat{y}_i$ and the GT map y_i, $\mathcal{L}_{seg}$ is obtained. By integrating $\mathcal{L}_{seg}$ with $\mathcal{L}_{recon}$, the model is trained to be robust under LACT conditions while preserving lesion-related details for accurate diagnosis.

2.3 Contrastive Abnormal Feature Loss

In medical imaging, subtle lesions or abnormalities are common and often challenging to detect. To identify subtle differences, anomaly detection typically transforms abnormal data into normal data and detects abnormalities by comparing the differences. However, this approach often introduces unintended changes in normal regions, leading to potential inaccuracies in lesion detection.

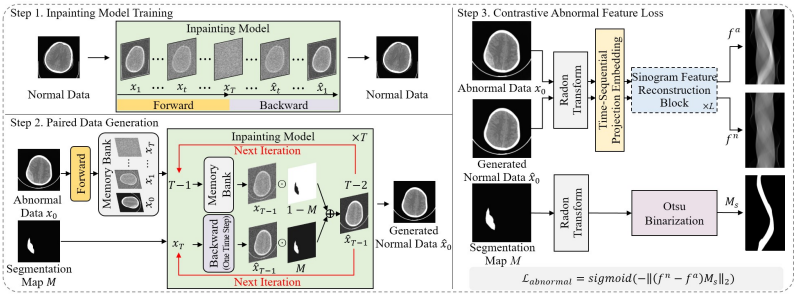

Fig. 3. Illustration of the contrastive abnormal feature loss.

Instead, we introduce an augmentation technique combining anomaly detection with inpainting to focus on regions of interest for generating paired data. Further, we propose a loss function that maximizes the feature differences between abnormal and normal regions, using the paired data to distinguish abnormalities.

As depicted in Step 1 of Fig. 3, similar to existing anomaly detection, the diffusion model is pre-trained on normal data. Building on this, in Step 2, we propose an inpainting method inspired by RePaint [15], where only the lesion regions are masked and transformed into normal representations as follows:

$$\hat{x}_{t-1} = \hat{x}_{t-1} \odot M + x_{t-1} \odot (1 - M), \tag{2}$$

where $\hat{x}_{t-1}$ denotes the reverse diffusion result estimated from $\hat{x}_t$ at time step t, x_{t-1} is the forward diffusion result of abnormal data x_0, M is the GT binary segmentation map, and $\odot$ is element-wise multiplication. The diffusion process is repeated over T time steps (set to 1000 [15]), preserving normal regions while modifying the masked abnormal regions to generate the paired normal data $\hat{x}_0$.

Using these paired abnormal and normal data obtained via this process, in Step 3, we apply the contrastive abnormal feature loss $\mathcal{L}_{abnorm}$ as follows:

$$\mathcal{L}_{abnorm} = sigmoid(-\|(f^a - f^n) \odot M_s\|_2), \tag{3}$$

where f^a and f^n denote the sinogram features obtained by the sinogram feature reconstruction using abnormal and normal data. The M_s denotes the abnormal region map, which is binarized using Otsu's method [17] from the sinogram transformed by the segmentation map. This loss maximizes the differences between normal and abnormal regions in the feature space for improved segmentation.

Finally, the total loss is defined using $\mathcal{L}_{seg}$ and $\mathcal{L}_{recon}$ in Section 2.2 as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{seg} + \mathcal{L}_{recon} + \mathcal{L}_{abnorm}. \tag{4}$$

3 Experiments

3.1 Experimental Setup

Dataset. The CTICH dataset [9] comprises CT scans from 82 patients with traumatic brain injuries, including 36 patients with ICH. It contains 2814 slices

Table 1. Quantitative comparison of medical image segmentation models on CTICH and FUMPE datasets across different angular ranges (360°, 120°, 90°, 60°, and 30°).

Segmentation Method	360°		120°		90°		60°		30°	
	DSC ↑	HD95 ↓	DSC ↑	HD95 ↓	DSC ↑	HD95 ↓	DSC ↑	HD95 ↓	DSC ↑	HD95 ↓
CTICH Dataset [9]										
TransUNet [3]	67.41	21.62	62.74	22.69	59.62	25.99	50.53	31.25	43.59	49.92
Swin-UNet [2]	61.92	36.33	54.64	37.50	51.44	40.36	43.31	42.72	32.47	64.32
TransCASCADE [20]	63.70	21.92	58.86	34.92	56.19	33.04	50.89	42.57	40.85	55.18
SCUNet++ [4]	65.70	26.29	58.95	21.75	57.44	22.93	50.45	28.34	41.94	47.63
SAMIHS [25]	68.12	20.40	61.36	21.50	57.34	22.71	52.19	34.06	43.75	38.65
Ours	**68.27**	**17.74**	**65.34**	**17.63**	**62.19**	**18.45**	**58.97**	**20.64**	**56.71**	**28.05**
FUMPE Dataset [16]										
TransUNet [3]	76.48	11.12	73.92	12.91	70.89	12.76	63.68	16.44	48.90	20.30
Swin-UNet [2]	72.31	20.05	69.97	21.73	65.54	21.88	61.43	26.38	44.84	31.64
TransCASCADE [20]	73.75	12.72	72.39	13.15	70.45	15.29	64.51	18.15	53.63	21.30
SCUNet++ [4]	76.39	13.76	73.61	14.59	69.43	17.24	61.73	19.65	46.75	26.64
SAMIHS [25]	77.04	12.14	74.33	12.59	71.92	12.81	66.64	15.33	55.50	23.60
Ours	**77.09**	**10.91**	**75.16**	**11.62**	**74.73**	**11.91**	**70.96**	**12.89**	**66.14**	**16.85**

Table 2. Quantitative comparison of medical image segmentation models with LACT reconstruction models on CTICH and FUMPE datasets across different angular ranges.

Reconstruction Method	Segmentation Method	120°		90°		60°		30°		Latency (s) ↓	Params (M) ↓	FLOPs (T) ↓
		DSC ↑	HD95 ↓	DSC ↑	HD95 ↓	DSC ↑	HD95 ↓	DSC ↑	HD95 ↓			
CTICH Dataset [9]												
DOLCE [12]	TransUNet [3]	61.15	28.21	59.99	20.03	56.65	21.80	51.62	32.45	8.88	379.59	57.19
	SAMIHS [25]	62.13	17.95	59.48	19.99	55.13	22.06	52.90	30.59	8.87	364.09	57.43
TIFADiffusion [26]	TransUNet [3]	64.64	18.87	59.78	19.74	56.78	25.85	49.65	34.74	70.83	169.34	530.22
	SAMIHS [25]	63.72	21.56	61.25	19.25	52.65	28.20	46.31	29.71	70.82	**153.92**	530.46
Ours		**65.34**	**17.63**	**62.19**	**18.45**	**58.97**	**20.64**	**56.71**	**28.05**	**0.45**	222.07	**0.92**
FUMPE Dataset [16]												
DOLCE [12]	TransUNet [3]	74.02	12.68	72.11	12.50	64.53	16.71	57.63	18.47	8.88	379.59	57.19
	SAMIHS [25]	74.65	12.71	72.76	12.15	68.14	13.55	59.67	20.72	8.87	364.09	57.43
TIFADiffusion [26]	TransUNet [3]	74.69	12.85	72.56	12.41	65.85	14.74	55.13	19.63	70.83	169.34	530.22
	SAMIHS [25]	74.95	12.18	73.59	12.33	68.49	14.01	60.57	17.38	70.82	**153.92**	530.46
Ours		**75.16**	**11.62**	**74.73**	**11.91**	**70.96**	**12.89**	**66.14**	**16.85**	**0.45**	222.07	**0.92**

at 512×512 pixels. The FUMPE dataset [16] comprises CT scans from 35 pulmonary embolism (PE) patients. It contains 8792 slices at 512×512 pixels. The labeled data in each dataset were split into training and testing sets with an 8:2 ratio. We simulated forward and back projection using fan-beam geometry with TorchRadon [21] across angular ranges of 120°, 90°, 60°, and 30°. We used DSC and 95% Hausdorff distance (HD95) for evaluation.

Implementation Details. This paper sets the number of detectors D to 768 and the total number of sinogram projections to 720 [14]. It sets N to 32, and L to 8. The experiments were conducted on an NVIDIA RTX 3090 GPU. TransSino used the Adam optimizer with learning rate 1e-3 for 150 epochs.

3.2 Experimental Results

Main Results. Table 1 evaluates the performance of medical image segmentation models on CTICH and FUMPE datasets across different angular ranges. All models were trained across all angular ranges at once. The results show that TransSino outperforms existing image domain models, even at 360° range, indicating the effectiveness of our reconstruction modules and losses.

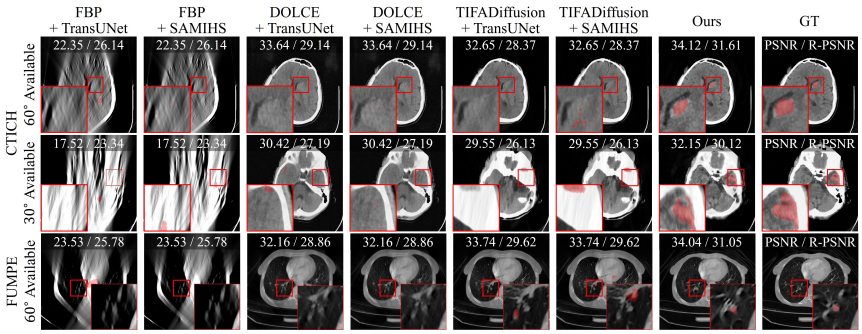

Fig. 4. Comparison of segmentation and reconstruction results on CTICH and FUMPE datasets. The values indicate average PSNR (dB) and R-PSNR (dB), respectively.

Table 3. (a) Ablation study of TransSino and (b, c) analysis of hyperparameters N and L on the CTICH dataset. DSC is measured as the average across all angular ranges.

Sequence Transformer	Prior Pattern Transformer	Auxiliary Reconstruction Loss	Contrastive Abnormal Feature Loss	DSC ↑
✓	✓	✓	✓	**60.80**
	✓	✓	✓	59.17
✓		✓	✓	57.73
✓	✓		✓	57.96
✓	✓	✓		59.50
				56.78

(a)

N	DSC ↑
16	60.76
32	**60.80**
64	58.16
128	55.33

(b)

L	DSC ↑
4	57.27
6	59.12
8	**60.80**
10	60.64

(c)

Table 2 evaluates the segmentation models combined with LACT reconstruction models on CTICH and FUMPE datasets across different angular ranges. The latency, Params, and FLOPs are presented as the sum of both models. The results indicate that TransSino consistently outperforms existing models, even when combined with image domain reconstruction models. Additionally, TransSino achieves much lower latency and FLOPs while maintaining comparable Params compared to the combined models. This performance gain is attributed to our approach using sinogram characteristics and auxiliary learning.

In Fig. 4, we compare the visual results of the reconstruction and segmentation models on CTICH and FUMPE datasets. Further, we evaluate the reconstruction performance averaged over each dataset in terms of PSNR and region PSNR (R-PSNR), measured within the lesion region using the GT map. The results indicate that TransSino provides more consistent outputs and better reconstructs lesion-related regions, such as ICH and PE, than the existing models.

Ablation Study. Table 3(a) presents the ablation study assessing the each component's contribution to TransSino. The results show that the first row, where all modules are activated (✓), achieves the highest DSC compared to the remaining rows, where individual modules are removed. This indicates that each module contributes to the performance gain. Further, Table 3(b and c) determine the optimal values for N and L, confirming best performance at 32 and 8.

4 Conclusion

This study presents an effective medical image segmentation strategy, surpassing existing methods under LACT conditions. TransSino predicts missing components by analyzing inherent order dependencies in the sinogram and disentangling sinusoidal patterns using prior knowledge for precise restoration. Further, it effectively distinguishes lesion regions via a unified loss function, ensuring robust performance while reducing complexity for diagnostic applications. Experimental results on ICH and PE in the brain and lungs show promising performance, suggesting generalizability to different anatomical regions. Although we use simulated LACT data by benchmarking existing methods [11,12,18,23,26,28], future work will focus on validation with real LACT data. Furthermore, we plan to extend our approach to 3D segmentation and reconstruction to better reflect its applicability to real-world volumetric data.

Acknowledgments. This work was supported by the IITP grant funded by the Korea government (MSIT) (No.2021-0-02068, RS-2022-00156287, RS-2023-00256629, RS-2024-00437718).

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

References

1. Anirudh, R., Kim, H., Thiagarajan, J. J., Mohan, K. A., Champley, K., & Bremer, T. (2018). Lose the views: Limited angle CT reconstruction via implicit sinogram completion. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 6343-6352).
2. Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., & Wang, M. (2022, October). Swin-unet: Unet-like pure transformer for medical image segmentation. In European conference on computer vision (pp. 205-218).
3. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., ... & Zhou, Y. (2021). Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306.
4. Chen, Y., Zou, B., Guo, Z., Huang, Y., Huang, Y., Qin, F., ... & Wang, C. (2024). SCUNet++: Swin-UNet and CNN Bottleneck Hybrid Architecture with Multi-Fusion Dense Skip Connection for Pulmonary Embolism CT Image Segmentation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (pp. 7759-7767).
5. Cho, J. H., & Fessler, J. A. (2013, October). Motion-compensated image reconstruction for cardiac CT with sinogram-based motion estimation. In 2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC) (pp. 1-5).
6. Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., ... & Ganapathy, R. (2024). The llama 3 herd of models. arXiv preprint arXiv:2407.21783.
7. Frysch, R., & Rose, G. (2015). Rigid motion compensation in interventional C-arm CT using consistency measure on projection data. In Medical Image Computing and Computer-Assisted Intervention. (pp. 298-306). Springer.

8. Hejrati, B., Banerjee, S., Glide-Hurst, C., & Dong, M. (2024, October). Conditional diffusion model with spatial attention and latent embedding for medical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention. (pp. 202-212). Springer.
9. Hssayeni, M., Croock, M., Salman, A., Al-khafaji, H., Yahya, Z., & Ghoraani, B. (2020). Computed tomography images for intracranial hemorrhage detection and segmentation. Intracranial hemorrhage segmentation using a deep convolutional model. Data, 5(1), 14. Available at https://physionet.org/content/ct-ich/1.3.1/.
10. Hssayeni, M. D., Croock, M. S., Salman, A. D., Al-Khafaji, H. F., Yahya, Z. A., & Ghoraani, B. (2020). Intracranial hemorrhage segmentation using a deep convolutional model. Data, 5(1), 14.
11. Hu, D., Zhang, Y., Liu, J., Luo, S., & Chen, Y. (2022). DIOR: Deep iterative optimization-based residual-learning for limited-angle CT reconstruction. IEEE Transactions on Medical Imaging, 41(7), 1778-1790.
12. Liu, J., Anirudh, R., Thiagarajan, J. J., He, S., Mohan, K. A., Kamilov, U. S., & Kim, H. (2023). DOLCE: A model-based probabilistic diffusion framework for limited-angle ct reconstruction. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 10498-10508).
13. Lin, T., Chen, Z., Yan, Z., Yu, W., & Zheng, F. (2024, October). Stable diffusion segmentation for biomedical images with single-step reverse process. In International Conference on Medical Image Computing and Computer-Assisted Intervention. (pp. 656-666). Springer.
14. Li, Z., Ma, C., Chen, J., Zhang, J., & Shan, H. (2023). Learning to distill global representation for sparse-view CT. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 21196-21207).
15. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., & Van Gool, L. (2022). Repaint: Inpainting using denoising diffusion probabilistic models. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 11461-11471).
16. Masoudi, M., Pourreza, H. R., Saadatmand-Tarzjan, M., Eftekhari, N., Zargar, F. S., & Rad, M. P. (2018). A new dataset of computed-tomography angiography images for computer-aided detection of pulmonary embolism. Scientific data, 5(1), 1-9. Available at https://doi.org/10.6084/m9.figshare.c.4107803.v1.
17. Otsu, N. (1975). A threshold selection method from gray-level histograms. Automatica, 11(285-296), 23-27.
18. Pan, J., Yu, H., Gao, Z., Wang, S., Zhang, H., & Wu, W. (2024). Iterative residual optimization network for limited-angle tomographic reconstruction. IEEE Transactions on Image Processing.
19. Paranjape, J. N., Sikder, S., Vedula, S. S., & Patel, V. M. (2024, October). S-sam: Svd-based fine-tuning of segment anything model for medical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention. (pp. 720-730). Springer.
20. Rahman, M. M., & Marculescu, R. (2023). Medical image segmentation via cascaded attention decoding. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (pp. 6222-6231).
21. Ronchetti, M. (2020). Torchradon: Fast differentiable routines for computed tomography. arXiv preprint arXiv:2009.14788. Available at https://github.com/matteo-ronchetti/torch-radon.

22. Roy, S., Koehler, G., Ulrich, C., Baumgartner, M., Petersen, J., Isensee, F., ... & Maier-Hein, K. H. (2023, October). Mednext: transformer-driven scaling of convnets for medical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention. (pp. 405-415). Springer.
23. Wang, C., Zhang, H., Li, Q., Shang, K., Lyu, Y., Dong, B., & Zhou, S. K. (2021). Improving generalizability in limited-angle CT reconstruction with sinogram extrapolation. In Medical Image Computing and Computer Assisted Intervention. (pp. 86-96). Springer.
24. Wang, X., Chen, H., Wu, Z., & Zhu, W. (2024). Disentangled representation learning. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(12), 9677-9696.
25. Wang, Y., Chen, K., Yuan, W., Tang, Z., Meng, C., & Bai, X. (2024, May). SAMIHS: adaptation of segment anything model for intracranial hemorrhage segmentation. In 2024 IEEE International Symposium on Biomedical Imaging (ISBI) (pp. 1-5).
26. Wang, Y., Li, Z., & Wu, W. (2024). Time-reversion fast-sampling score-based model for limited-angle CT reconstruction. IEEE Transactions on Medical Imaging, 43(10), 3449-3460.
27. Zhang, J., Wan, T., MacDonald, M. E., Menon, B. K., Qiu, W., & Ganesh, A. (2024, October). Synchronous Image-Label Diffusion with Anisotropic Noise for Stroke Lesion Segmentation on Non-Contrast CT. In International Conference on Medical Image Computing and Computer-Assisted Intervention. (pp. 433-443). Springer.
28. Zhou, B., Zhou, S. K., Duncan, J. S., & Liu, C. (2021). Limited view tomographic reconstruction using a cascaded residual dense spatial-channel attention network with projection data fidelity layer. IEEE transactions on medical imaging, 40(7), 1792-1804.