



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# M<sup>4</sup>oE: A Foundation Model for Medical Multimodal Image Segmentation with Mixture of Experts

Yufeng Jiang<sup>1</sup>[0009-0004-4987-2683] and Yiqing Shen<sup>2</sup>(✉)[0000-0001-7866-3339]

<sup>1</sup> Department of Computer Science, Hong Kong Baptist University, Hong Kong SAR, China


<sup>2</sup> Department of Computer Science, Johns Hopkins University, Baltimore, USA  
yshen92@jhu.edu

**Abstract.** Medical imaging data is inherently heterogeneous across different modalities and clinical centers, posing unique challenges for developing generalizable foundation models. Conventional entails training distinct models per dataset or using a shared encoder with modality-specific decoders. However, these approaches incur heavy computational overheads and suffer from poor scalability. To address these limitations, we propose the Medical Multimodal Mixture of Experts (M<sup>4</sup>oE) framework, leveraging the SwinUNet architecture. Specifically, M<sup>4</sup>oE comprises modality-specific experts; each separately initialized to learn features encoding domain knowledge. Subsequently, a gating network is integrated during fine-tuning to modulate each expert’s contribution to the collective predictions dynamically. This enhances model interpretability and generalization ability while retaining expertise specialization. Simultaneously, the M<sup>4</sup>oE architecture amplifies the model’s parallel processing capabilities, and it also ensures the model’s adaptation to new modalities with ease. Experiments across three modalities reveal that M<sup>4</sup>oE can achieve 3.45% over STU-Net-L, 5.11% over MED3D, and 11.93% over SAM-Med2D across the MICCAI FLARE22, AMOS2022, and ATLAS2023 datasets. Moreover, M<sup>4</sup>oE showcases a significant reduction in training duration with 7 hours less while maintaining a parameter count that is only 30% of its compared methods. The code is available at <https://github.com/JefferyJiang-YF/M4oE>.

**Keywords:** Foundation Model · Multimodal Medical Image Segmentation · Mixture of Experts (MoE) · SwinUNet.

## 1 Introduction

The advent of vision foundation models such as Segment Anything Model (SAM) and Stable Diffusion [18] has catalyzed remarkable progress in natural image understanding [14]. However, translating these advancements to medical imaging poses notable challenges [7] due to the heterogeneity across modalities like MRI, CT, and X-rays, each exhibiting distinct data characteristics [8]. Consequently,




**Fig. 1.** Paradigms for multimodal medical segmentation: (a) One network for each modality; (b) A shared encoder with modality-specific decoders; (c) Union of labels with pseudo-label training; (d) Our proposed M<sup>4</sup>oE with shared encoder and decoder.

medical imaging necessitates specialized processing beyond the capabilities of models designed for more homogeneous natural images. Historically, coping with multimodal medical data heterogeneity entailed developing customized models targeting individual datasets, as depicted in Fig. 1(a). Although minimizing initial complexity, this approach introduces substantial overheads for adapting models to new datasets. Alternatively, following the multi-task learning design, some methods utilize a shared encoder with modality-specific decoders to extract common representations while retaining tailored processing [3,20], as in Fig. 1(b). However, computational requirements can still scale poorly as data diversity grows [11]. Unified encoder-decoder schemes also exist to standardize label processing [12], as shown in Fig. 1(c), but may compromise extensibility to new modalities [1].

To address these shortcomings, we introduce Medical Multimodal Mixture of Experts (M<sup>4</sup>oE), built upon SwinUNet [2], to effectively process multimodal medical images segmentation. The conventional Transformer architecture struggles to adapt to modality-specific attributes optimally [19]. To overcome this gap, M<sup>4</sup>oE initializes distinct experts focusing on unique aspects of each modality (Fig. 1(d)). A gating network then dynamically combines expert outputs during inference. The design of M<sup>4</sup>oE enhances parallel processing capabilities and simplifies adapting to new modalities without extensive reconfiguration.

The contribution of this work is three-fold, summarized as follows. 1) We propose an innovative foundation model for multimodal medical image segmentation based upon a mixture of experts, named M<sup>4</sup>oE. It allows efficient scaling to large heterogeneous datasets without drastic parameter increases, striking an optimal balance between broad applicability and computational efficiency. 2) We propose M<sup>4</sup>oE’s two-phase training with modality-specific experts and, a fusing gating network, and a shuffle-adaptable linear projection architecture for multimodality and label mapping. 3) Our comprehensive experimental evaluations



**Fig. 2.** Overall framework of M<sup>4</sup>oE implemented on a SwinUNet backbone. At the core is the M<sup>4</sup>oE with a gating network, which dynamically selects specialized experts for different imaging modalities (CT, CE-MRI, MRI); each expert learns representation for its respective data characteristics.

indicate that M<sup>4</sup>oE achieves competitive performance on multimodal medical datasets and shows promising transferability, with reduced need for reconfiguration across different modalities.

## 2 Methods

We present M<sup>4</sup>oE building on SwinUNet [2], where our model substitutes the MLP component with a Mixture-of-Experts designed specifically for multimodal medical image segmentation to improve representation.

**Swin Transformer Backbone** The core backbone of our architecture leverages Swin Transformer blocks [15]. Contrasting conventional multi-headed self-attention (MHSA), Swin Transformers employ a local, sliding-window variant termed window-based MHSA (W-MSA). As illustrated in Fig. 2, each Swin Transformer block comprises a LayerNorm, a Multilayer Perceptron (MLP), W-MSA, and a shifted-window MHSA (SW-MSA). The sliding window approach

limits self-attention computation to local patch regions, reducing complexity compared to global MHSA.

**Medical Multi-Modal Mixture-of-Experts (M<sup>4</sup>oE) Formulation** Our key innovation involves substituting the MLP layer in each Swin Transformer block with a M<sup>4</sup>oE, as depicted in Fig. 1(d). This comprises a gating network along with a set of modality-specific expert networks. The experts themselves consist of two-layer MLPs, utilizing GELU for non-linear activation. Each expert network specializes in learning representations corresponding to a particular imaging modality. This allows dynamically allocating computational resources conditional on data complexity, enhancing the model’s representation capacity. The gating network leverages a simple single-layer MLP activated by a soft-max function to weigh the contribution of each expert dynamically based on the input.

**M<sup>4</sup>oE in Swin Transformer** We incorporate the proposed M<sup>4</sup>oE into Swin Transformer blocks as follows:

$$\begin{aligned}\tilde{z}^l &= \text{W-MSA}(\text{LN}(z^{l-1})) + z^{l-1}, \\ z^l &= \text{M}^4\text{oE}(\text{LN}(\tilde{z}^l)) + \tilde{z}^l, \\ \tilde{z}^{l+1} &= \text{SW-MSA}(\text{LN}(z^l)) + z^l, \\ z^{l+1} &= \text{M}^4\text{oE}(\text{LN}(\tilde{z}^{l+1})) + \tilde{z}^{l+1},\end{aligned}\tag{1}$$

where  $\tilde{z}$  and  $z$  denote the W/SW-MSA and M<sup>4</sup>oE representations of the  $l^{\text{th}}$  block respectively. Within M<sup>4</sup>oE, the gating network assigns weights to each expert’s output as follows:

$$\begin{aligned}\text{weight}_G &= \text{SoftMax}(W_g \cdot \mathbf{x} + \mathbf{b}_g) = \frac{\exp(\text{weights}_G)}{\sum_j \exp(\text{weights}_{G_j})}, \\ \text{where } \text{expert}_i(\mathbf{x}) &= W_{o,i} \cdot \text{GELU}(W_{h,i} \cdot \mathbf{x} + \mathbf{b}_{h,i}) + \mathbf{b}_{o,i},\end{aligned}\tag{2}$$

where  $\text{weight}_G$  is the gating weights to the input  $\mathbf{x}$ , the gating weights are computed using the SoftMax function for assigning different weights based on other representations.  $W_{h,i}$ ,  $\mathbf{b}_{h,i}$ ,  $W_{o,i}$ , and  $\mathbf{b}_{o,i}$  are the learnable weights and biases of the  $i^{\text{th}}$  expert network. The outputs of all experts are computed and stacked together to form a tensor *i.e.*,  $\text{expert}_1(\mathbf{x}), \text{expert}_2(\mathbf{x}) \dots \text{expert}_{\text{num\_experts}}(\mathbf{x})$ . It follows that the output of the M<sup>4</sup>oE module is the weighted sum of all expert outputs, namely

$$\text{outputs} = \sum_{i=1}^n \text{weight}_i \cdot \text{expert}_i(\mathbf{x}),\tag{3}$$

where  $\text{weight}_i$  is the  $i^{\text{th}}$  element of the weights vector, representing the weight assigned to the  $i^{\text{th}}$  expert’s output.

**Table 1.** Comparison of M<sup>4</sup>oE against prior arts in terms of crucial model design choices and computational requirements, including pre-training strategy, input dimensionality, multimodality handling capability, number of parameters, training duration, and parameters added when introducing new modalities.


Methods	Pre-train	2D/3D Multimodality	#Param	Training hours	Introduce new modality
DoDNet [20]	✗	3D ✓	17.3M	48 hours	-
MED3D [3]	✗	3D ✓	117.51M	16 hours	21M
Liu’s [13]	✗	3D ✗	6.66M	19.5 hours	6.66M
STU-Net-L [9]	✓ Supervised	3D ✗	440.30M	156.8 hours	440.30M
Ours	✓ Self-supervised	2D ✓	27.31M	12 hours	10M

**Dynamic Selection for Linear Projection Head** A core challenge in multimodal segmentation involves handling varying class numbers across modalities. We introduce a dynamic selection mechanism through customizable linear projection heads inserted before output to address this. First, feature maps are padded to equalize dimensions across all samples regardless of modality. Subsequently, each sample’s linear projection layer dynamically selects only the logits corresponding to the classes present. Thereby, the final output dimensions are aligned with the native label distribution in a modular, input-conditional manner. This process grants invariance to shuffling across modalities during training. Meanwhile, the lightweight projections minimize overheads, allowing efficient re-configuration of new datasets.

**Two-phase Training Strategy** We propose a novel two-phase training strategy consisting of expert pre-training and gate network fine-tuning phases. In the expert pre-training phase, we leverage the Masked Autoencoder(MAE) [6] to pre-train each expert. This enables it to learn meaningful representations for its modality without the gating network. Subsequently, in the gate network fine-tuning phase, we load the pre-trained encoder and expert parameters and incorporate the gating network. We then fine-tune the entire M<sup>4</sup>oE model end-to-end on downstream segmentation tasks. This two-phase approach allows the experts to specialize in their respective modalities before learning to selectively combine their representations via the gating network.

### 3 Experiments

**Implementations** All experiments leveraged a single RTX 4090 GPU system with 24GB memory, running PyTorch 1.11.0, Python 3.8 on Ubuntu 20.04, and CUDA 11.3. To initialize feature representations, the M<sup>4</sup>oE in encoder was pre-trained via MAE [5,6]. The decoder, optimization settings, input resolution, and patch size were configured identically to SwinUNet [2] for fair comparison. Specifically, training employed a batch size of 36 and weight decay of  $10^{-4}$  using AdamW optimization.



**Fig. 3.** Visualization of segmentation results from various methods. The first and second rows depict FLARE22 with CT images, the third row showcases ATLAS2023 with CE-MRI images, the fourth and sixth rows display AMOS2022 with CT images, and the fifth row presents MRI images from AMOS2022. The six columns from left to right correspond to the original image, the ground truth (GT), the DoDNet results, the MED3D results, the SAM-MED2D results, and our  $M^4oE$  results.

**Datasets** We leverage three multimodal medical image segmentation datasets to evaluate  $M^4oE$ 's performance, namely FLARE22 [16], AMOS22 [10], and ATLAS23 [17]. Specifically, from FLARE22, we utilize 50 labeled CT scan cases. AMOS22 provides 600 labeled cases comprising 500 CT and 100 MRI abdominal scans. Additionally, we incorporate 60 cases of contrast-enhanced MRI (CE-MRI) liver scans with labels from ATLAS23. We follow the previous work [9] for data split. While FLARE22 and AMOS22 focus on segmenting abdominal organs, ATLAS23 specializes in segmenting liver and tumors. This allows assessing  $M^4oE$ 's capability in handling multimodal datasets across diverse anatomical segmentation tasks.

**Results** As shown in Table 1, DoDNet [20], MED3D [3], and Liu's [13] methods lack pre-training, potentially limiting learned representations. Although focusing on 3D data, DoDNet and MED3D partially mitigate this via multimodality. MED3D's efficiency is evident with its short training duration despite a high

**Table 2.** Comparison of DSC (%) across multimodal segmentation methods. Our M<sup>4</sup>oE model outperforms others, achieving the highest DSC in FALRE and AMOS-CT tasks. It shows adaptability and improvement (\*) upon introducing a new modality, contrary to DoDNet and MED3D, which see performance drops. Unlike STU-Net-L’s single-dataset focus, our model’s robustness is attributed to training on three varied datasets. SAM-Med2D’s performance is evaluated using a single data point. Results are highlighted in red for best and blue for second best.

Datasets	Method	Liver	Left Kidney	Right Kidney	Gallbladder	Pancreas	Spleen	Stomach	Mean
FLARE22	DoDNet [20]	96.64	95.43	95.32	74.58	79.48	90.91	90.85	89.03
	MED3D [3]	90.36	87.23	88.47	67.84	73.82	91.98	92.30	84.53
	Liu’s [13]	89.30	76.32	78.40	55.35	58.86	83.07	67.82	72.73
	STU-Net-L [9]	94.48	87.51	86.68	76.31	82.00	91.27	89.31	86.79
	SAM-Med2D [4]	80.37	77.96	78.19	50.37	53.12	82.47	75.01	71.07
	Ours	97.08	94.62	94.92	79.38	76.99	95.74	89.78	89.78
AMOS-CT	DoDNet [20]	92.14	83.93	82.48	70.73	73.69	90.58	85.14	82.67
	MED3D [3]	88.25	85.11	85.29	65.12	59.78	88.37	81.10	79.00
	STU-Net-L [9]	95.47	88.37	89.74	76.92	82.64	93.24	87.03	87.63
	SAM-Med2D [4]	67.51	71.09	68.31	57.93	51.24	76.12	72.93	66.45
	Ours	95.59	88.72	88.31	73.94	61.39	90.59	82.78	83.04
AMOS-CT+MRI*	DoDNet [20]	89.31	80.33	82.89	70.13	68.32	85.80	83.23	80.01
	MED3D [3]	83.42	82.11	86.33	69.12	63.20	86.23	80.79	78.74
	STU-Net-L [9]	95.59	89.00	90.10	79.97	94.08	94.37	94.42	91.08
	SAM-Med2D [4]	55.29	63.27	62.57	48.18	50.12	61.29	68.10	58.40
	Ours	94.38	89.20	88.16	76.37	69.12	91.23	85.34	84.83

parameter count, while Liu’s method, with fewer parameters, might be better generalized. Meanwhile, STU-Net-L’s [9] supervised pre-training potentially improves initial learning but necessitates labeled data, and its high parameter count and extended training time could be drawbacks. Our approach introduces self-supervised pre-training on 2D data, eliminating the need for labeled datasets and showing cost-effectiveness. It balances model complexity and resource efficiency, proving advantageous for practical applications where performance and efficiency are paramount. As shown in Table 2, our approach excels in the FLARE22 and AMOS-CT datasets, with the highest mean dice scores (DSC), indicating its suitability for multimodal tasks and accurate segmentation. Our method shows slight performance decreases when introducing MRI modality, evidencing its adaptability, whereas DoDNet and MED3D exhibit notable performance drops. Our model’s training across diverse datasets likely contributes to its robustness, as opposed to STU-Net-L, which although achieving high DSC in a single-task setting, lacks versatility. SAM-Med2D’s lower performance might stem from its 2D focus, potentially failing to capture 3D medical imaging complexities. Our method outperforms the ATLAS2023 dataset, maintaining robustness with new modalities and leading in DSC and mIoU metrics, as seen in Table 3. This underscores the limitations of DoDNet and MED3D in modality variations, while our model’s balanced approach offers a promising direction for multimodal medical image segmentation.

**Table 3.** Results on different methods on ATLAS2023, we use one point to test SAM-Med2D on test sets. The DoDNet and MED3D are not robust when the new modality is introduced. The text is in red and in blue for the best and second-best results.

Methods	DSC (%)			mIoU (%)		
	Liver	Tumor	Mean	Liver	Tumor	Mean
DoDNet[20]	79.54	24.03	52.34	75.31	19.3	47.31
MED3D [3]	73.12	19.55	46.34	65.3	16.82	41.06
SAM-Med2D [4]	80.02	51.28	65.65	77.34	48.88	63.11
Ours	89.43	57.53	73.48	86.34	55.32	70.83

**Table 4.** Ablation study on different model configurations. Encoder or decoder marked with  $\times$  indicates the use of an MLP as the output layer for the Swin Transformer block. In contrast,  $\checkmark$  denotes that the encoder or decoder incorporates our M<sup>4</sup>oE structure. 'GN' stands for gating network. The ablation studies reveal that employing the M<sup>4</sup>oE architecture in both encoder and decoder significantly enhances model performance. The text is in red and in blue for the best and second-best results.

Method				DSC (%)								
Encoder	Decoder	GN	M <sup>4</sup> oE	Liver	Liver Tumor	Left Kidney	Right Kidney	Gallbladder	Pancreas	Spleen	Stomach	Mean
$\times$	$\times$	$\times$	$\times$	93.61	51.23	84.65	84.51	69.34	63.21	92.37	81.20	77.52
$\checkmark$	$\times$	$\checkmark$	$\checkmark$	93.21	55.12	89.27	89.15	75.41	68.32	91.12	83.25	80.61
$\checkmark$	$\times$	$\times$	$\checkmark$	93.58	53.21	90.01	90.60	76.01	65.33	90.87	82.79	80.30
$\times$	$\checkmark$	$\checkmark$	$\checkmark$	94.08	53.25	88.98	89.01	75.76	69.31	92.37	85.81	81.07
$\times$	$\checkmark$	$\times$	$\checkmark$	94.30	53.19	90.12	88.91	74.11	68.45	91.36	85.43	80.73
$\checkmark$	$\checkmark$	$\times$	$\checkmark$	93.12	55.37	88.67	87.56	75.39	66.29	90.12	83.04	79.95
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	94.12	57.53	90.85	90.46	76.56	69.16	92.52	85.97	82.15

**Ablation Study** The ablation study in Table 4 assesses the impact of implementing the M<sup>4</sup>oE architecture in different components of a segmentation network. It indicates that integrating the M<sup>4</sup>oE structure into both the encoder and decoder significantly enhances the model’s segmentation performance, as evidenced by the highest mean DSC score of 82.15%. Models using M<sup>4</sup>oE in either the encoder or decoder showed improved results over the baseline model without M<sup>4</sup>oE, but the full incorporation into both yielded the best outcomes. The M<sup>4</sup>oE architecture benefits the backbone model by enabling effective multimodal data integration, capturing diverse contextual information at multiple scales, and modulating features to emphasize relevant patterns for accurate segmentation. Additionally, the synergy between the encoder and decoder when both use the M<sup>4</sup>oE structure results in more coherent feature learning and reconstruction. The presence of a gating network also contributes to performance improvements, suggesting its role in enhancing feature flow control within the network.

## 4 Conclusion

The proposed Medical Multimodal Mixture of Experts (M<sup>4</sup>oE) presents a novel and efficient direction toward developing foundation models for multimodal medical image segmentation. By effectively integrating modality-specific experts



within a flexible gating mechanism, M<sup>4</sup>oE demonstrates accurate and adaptable segmentation across diverse imaging datasets. We believe the expert specialization and dynamic ensemble concepts presented here will provide guiding principles for tackling multimodal challenges in medical imaging and beyond. For future direction, evaluating M<sup>4</sup>oE over larger-scale datasets, augmenting with uncertainty estimation, and enhancing model transparency could further facilitate clinical integration.

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

## References

1. Anwar, S.M., Majid, M., Qayyum, A., Awais, M., Alnowami, M., Khan, M.K.: Medical image analysis using convolutional neural networks: a review. *Journal of medical systems* **42**, 1–13 (2018)
2. Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., Wang, M.: Swin-unet: Unet-like pure transformer for medical image segmentation. In: *European conference on computer vision*. pp. 205–218. Springer (2022)
3. Chen, S., Ma, K., Zheng, Y.: Med3d: Transfer learning for 3d medical image analysis. *arXiv preprint arXiv:1904.00625* (2019)
4. Cheng, J., Ye, J., Deng, Z., Chen, J., Li, T., Wang, H., Su, Y., Huang, Z., Chen, J., Jiang, L., et al.: Sam-med2d. *arXiv preprint arXiv:2308.16184* (2023)
5. Dai, Y., Liu, F., Chen, W., Liu, Y., Shi, L., Liu, S., Zhou, Y., et al.: Swin mae: Masked autoencoders for small datasets. *Computers in Biology and Medicine* **161**, 107037 (2023)
6. 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)
7. Heiliger, L., Sekuboyina, A., Menze, B.H., Egger, J., Kleesiek, J.: Beyond medical imaging - a review of multimodal deep learning in radiology (2022). <https://doi.org/10.36227/techrxiv.19103432.v1>
8. Huang, H., Lin, L., Tong, R., Hu, H., Zhang, Q., Iwamoto, Y., Han, X., Chen, Y.W., Wu, J.: Unet 3+: A full-scale connected unet for medical image segmentation. In: *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. pp. 1055–1059. IEEE (2020)
9. Huang, Z., Wang, H., Deng, Z., Ye, J., Su, Y., Sun, H., He, J., Gu, Y., Gu, L., Zhang, S., et al.: Stu-net: Scalable and transferable medical image segmentation models empowered by large-scale supervised pre-training. *arXiv preprint arXiv:2304.06716* (2023)
10. Ji, Y., Bai, H., Ge, C., Yang, J., Zhu, Y., Zhang, R., Li, Z., Zhang, L., Ma, W., Wan, X., Luo, P.: AMOS: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (eds.) *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022* (2022), [http://papers.nips.cc/paper\\_files/paper/2022/hash/ee604e1bedbd069d9fc9328b7b9584be-Abstract-Datasets\\_and\\_Benchmarks.html](http://papers.nips.cc/paper_files/paper/2022/hash/ee604e1bedbd069d9fc9328b7b9584be-Abstract-Datasets_and_Benchmarks.html)

11. Johnson, A.E., Pollard, T.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Peng, Y., Lu, Z., Mark, R.G., Berkowitz, S.J., Horng, S.: Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. arXiv preprint arXiv:1901.07042 (2019)
12. Liu, P., Deng, Y., Wang, C., Hui, Y., Li, Q., Li, J., Luo, S., Sun, M., Quan, Q., Yang, S., et al.: Universal segmentation of 33 anatomies. arXiv preprint arXiv:2203.02098 (2022)
13. Liu, W., Xu, W., Yan, S., Wang, L., Li, H., Yang, H.: Combining self-training and hybrid architecture for semi-supervised abdominal organ segmentation. FLARE@MICCAI (2022). [https://doi.org/10.1007/978-3-031-23911-3\\_25](https://doi.org/10.1007/978-3-031-23911-3_25), <https://arxiv.org/abs/2207.11512v4>
14. Liu, X., Zhou, T., Wang, Y., Wang, Y., Cao, Q., Du, W., Yang, Y., He, J., Qiao, Y., Shen, Y.: Towards the unification of generative and discriminative visual foundation model: A survey. arXiv preprint arXiv:2312.10163 (2023)
15. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
16. Ma, J., Zhang, Y., Gu, S., Ge, C., Ma, S., Young, A., Zhu, C., Meng, K., Yang, X., Huang, Z., Zhang, F., Liu, W., Pan, Y., Huang, S., Wang, J., Sun, M., Xu, W., Jia, D., Choi, J.W., Alves, N., de Wilde, B., Koehler, G., Wu, Y., Wiesenfarth, M., Zhu, Q., Dong, G., He, J., the FLARE Challenge Consortium, Wang, B.: Unleashing the strengths of unlabeled data in pan-cancer abdominal organ quantification: the flare22 challenge. arXiv preprint arXiv:2308.05862 (2023)
17. Quinton, F., Popoff, R., Presles, B., Leclerc, S., Meriaudeau, F., Nodari, G., Lopez, O., Pellegrinelli, J., Chevallier, O., Ginjac, D., Vrigneaud, J.M., Alberini, J.L.: A tumour and liver automatic segmentation ( atlas) dataset on contrast-enhanced magnetic resonance imaging for hepatocellular carcinoma. Data (2023). <https://doi.org/10.3390/data8050079>, <https://www.mdpi.com/2306-5729/8/5/79>
18. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
19. Shin, A., Ishii, M., Narihira, T.: Perspectives and prospects on transformer architecture for cross-modal tasks with language and vision. International Journal of Computer Vision (2021). <https://doi.org/10.1007/s11263-021-01547-8>, <https://arxiv.org/abs/2103.04037v2>
20. Zhang, J., Xie, Y., Xia, Y., Shen, C.: Dodnet: Learning to segment multi-organ and tumors from multiple partially labeled datasets. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1195–1204 (2021)