

Robust Semi-supervised Multimodal Medical Image Segmentation via Cross Modality Collaboration

Xiaogen Zhou¹, Yiyun Sun², Min Deng²,
Winnie Chiu Wing Chu²(✉), and Qi Dou¹(✉)

¹Department of Computer Science and Engineering,
The Chinese University of Hong Kong, Hong Kong, China

² Department of Imaging and Interventional Radiology,
The Chinese University of Hong Kong, Hong Kong, China

Abstract. Multimodal learning leverages complementary information derived from different modalities, thereby enhancing performance in medical image segmentation. However, prevailing multimodal learning methods heavily rely on extensive well-annotated data from various modalities to achieve accurate segmentation performance. This dependence often poses a challenge in clinical settings due to limited availability of such data. Moreover, the inherent anatomical misalignment between different imaging modalities further complicates the endeavor to enhance segmentation performance. To address this problem, we propose a novel semi-supervised multimodal segmentation framework that is robust to scarce labeled data and misaligned modalities. Our framework employs a novel cross modality collaboration strategy to distill modality-independent knowledge, which is inherently associated with each modality, and integrates this information into a unified fusion layer for feature amalgamation. With a channel-wise semantic consistency loss, our framework ensures alignment of modality-independent information from a feature-wise perspective across modalities, thereby fortifying it against misalignments in multimodal scenarios. Furthermore, our framework effectively integrates contrastive consistent learning to regulate anatomical structures, facilitating anatomical-wise prediction alignment on unlabeled data in semi-supervised segmentation tasks. Our method achieves competitive performance compared to other multimodal methods across three tasks: cardiac, abdominal multi-organ, and thyroid-associated orbitopathy segmentations. It also demonstrates outstanding robustness in scenarios involving scarce labeled data and misaligned modalities. Code is available at: <https://github.com/med-air/CMC>.

Keywords: Semi-supervised learning · Multimodal image segmentation

1 Introduction

Multimodal learning is a long-standing topic with great importance for medical image analysis. It leverages synergistic information of multiple modalities

to outperform single-modality solutions. Many multimodal methods have been developed for medical image segmentation. These methods either concatenate a large array of well-annotated multiple Magnetic Resonance Imaging (MRI) and computed tomography (CT) as input [14,15], or fuse higher-level information from pre-aligned modalities within a latent space [2,8]. However, in real-world clinical settings, the availability of extensive collections of meticulously annotated and anatomically aligned datasets from various modalities cannot be always guaranteed. This is primarily due to the expensive costs associated with detailed annotation and a shortage of experienced clinicians available for labeling tasks. Consequently, robustness to sparsely labeled data or misaligned modalities during inference is essential for a widely-applicable multimodal learning method.

A typical solution to address the scarcity of labeled data is to employ semi-supervised learning techniques [1,16]. These methods can effectively leverage the limited labeled data in conjunction with abundant unlabeled data to train deep learning models, thereby mitigating the dependence on costly labeled data while maintaining satisfactory performance. However, most existing semi-supervised learning methods are inherently constrained to single-modality input, failing to capitalize on the wealth of multimodal data frequently encountered in daily clinical routines. Furthermore, Zhu et al. [18] and Zhang et al. [13] have attempted to address this by applying semi-supervised learning to minimize consistency loss on abundant unlabeled data across different modalities. However, different modalities vary in intensity distributions and appearances, and simply minimizing consistency loss may not be sufficient for effective feature fusion and alignment. To address this challenge, a recent work [2] introduces a modality-collaborative semi-supervised framework. By leveraging modality-specific knowledge, this method endeavors to enhance multimodal image segmentation, thereby facilitating more effective feature fusion and alignment.

For the effective extraction of modality-invariant representations that capture critical information and enable feature fusion across various modalities, it is beneficial to identify and leverage modality-specific knowledge. This can be accomplished through modality-independent feature extraction, which distills the characteristics and semantic representations of each modality [2,3]. Additionally, integrating a modality-collaborative strategy can significantly enhance feature fusion and alignment across different modalities, ensuring effective synergistic integration of multimodal data [2,14]. However, these works fail to enforce constraints to regularize the underlying anatomical structures during the feature fusion stage, which is crucial for maintaining alignment.

In this paper, we propose a novel semi-supervised multimodal framework that incorporates cross modality collaboration and contrastive consistent learning. This framework is robust to scenarios with limited labeled data and misaligned modalities. Our framework exploits a novel cross modality-collaboration strategy to regularize the underlying channel-wise features and distill modality-independent knowledge. The modality-independent knowledge extracted from each modality is fused to a common fusion layer, which contains collaborative information for segmentation. To effectively extract modality-independent in-




Fig. 1. The pipeline of our framework. It consists of three main components. (1) Two 3D foundation model-driven encoders are employed to capture 3D volumetric representations from different modalities. (2) A novel cross-modality collaborative strategy is introduced to fuse information and align feature maps. (3) A contrastive consistent learning module is exploited to generate consistent predictions for unlabeled data.

formation, we design modality-specific encoders to capture the distinctive semantic features intrinsic to each modality. Moreover, we introduce a novel contrastive consistent learning module to mitigate inconsistencies and align anatomical structures between prediction maps from unlabeled data across different modalities, thereby facilitating contrastive consistent learning process among image modalities. We validate our proposed method on the task of 3D semi-supervised multimodal medical image segmentation across three datasets for different tasks. The results demonstrate that our method achieves competitive performance compared to state-of-the-art methods for misaligned modality data and exhibits high robustness in situations with scarce labeled data.

2 Methodology

An overview of our framework is presented in Fig. 1. We first employ two 3D modality-specific encoders to get the initial modality-independent features from multimodal inputs. Then, we introduce a cross-modality collaboration strategy to facilitate feature fusion and channel-wise alignment. Subsequently, we introduce a contrastive consistent learning module aimed at reducing inconsistencies in anatomical structures across prediction maps derived from the unlabeled data. The details of our method are elaborated in the following sections.

2.1 Cross Modality Collaboration for Alignment and Feature Fusion

We denote the multimodal images by $\mathcal{D}^l = \{(x_{l,i}^a, y_i^a)\}_{i=1}^N, \{(x_{l,i}^b, y_i^b)\}_{i=1}^N\}$ and $\mathcal{D}^u = \{(x_{u,j}^a)\}_{j=1}^M, \{(x_{u,j}^b)\}_{j=1}^M\}$, where $x_{l,i}$ and $x_{u,j}$ denote the labeled and unlabeled data in modality a and b, respectively. The y_i^a and y_i^b are the corresponding segmentation masks for $x_{l,i}^a$ and $x_{l,i}^b$. The N and M denote the number of labeled and unlabeled samples, and $N \ll M$. Each modality x_i^a and x_i^b is input

to modality-specific encoder E^a and E^b , respectively, which are fine-tuned from the pre-trained encoder of SAM-Med3D [10]. Moreover, as shown in Fig. 1 (b), we integrate an adapter module within each Vision Transformer (ViT) block. This adapter serves as a bottleneck infrastructure and consists of a sequence of two 3D convolutional layers with different kernel sizes.

Then, we obtain its distinctive semantic (DS) features by $F_{ds}^a = E^a(x_i^a)$ and $F_{ds}^b = E^b(x_i^b)$. For effective extraction of distinctive semantic features F_{ds}^a and F_{ds}^b , which are expected to regularize underlying anatomical structures from a perspective of channel-wise feature alignment across modalities, we introduce a Channel-wise Semantic Consistency (CSC) loss which is used to align the channel-wise features. In specific, the $\mathcal{L}_{CSC}$ loss can be defined as follows:

$$\mathcal{L}_{CSC} = \min_{E^a, E^b} \sum_c -\log \frac{\exp(\text{Sim}((F_{ds}^a)_c, (F_{ds}^b)_c))}{\sum_c \exp(\text{Sim}((F_{ds}^a)_c, (F_{ds}^b)_c))}, \quad (1)$$

where $c \in \{1, 2, \dots, C\}$ denotes the c_{th} channel of two features, i.e., F_{ds}^a and F_{ds}^b , respectively. $\text{Sim}(\cdot, \cdot)$ represents the cosine similarity [9] used to measure the channel-wise semantic feature similarity between the latent representations.

After the feature-wise alignment procedure, which mitigates the influence of modality-specific appearance features from multimodal data, we introduce a novel Modality-Independent Awareness (MIA) module to further harness modality independent knowledge from each modality for effective feature fusion. As is known from literature of deep neural networks, modality-independent knowledge is useful in learning generalized and robust representations across various imaging modalities [2, 17]. As depicted in Fig. 1, the MIA module integrates modality-aware attention mechanisms [15], 3D convolutional layers, and a fusion layer. Subsequently, the features F_{mia}^a and F_{mia}^b generated by each MIA module are subsequently merged in a fusion layer to produce the fused feature.

2.2 Contrastive Consistent Learning with Contrastive Consistency

Contrastive consistency has shown effectiveness in semi-supervised learning [2, 13]. The formulation of the contrastive consistency term is based on the correlation and differential information across various modalities. The consistency of their predictions acts as the impetus for semi-supervised cross-modal knowledge reciprocal learning. This also holds true for our scenario. We first build two decoders D^a and D^b , which are fine-tuned the pre-trained foundation model decoder [6] on our datasets. During training, the weights of our model are optimized using a supervised loss function $\mathcal{L}_{Sup}$ on labeled data, which is defined as follows:

$$\mathcal{L}_{Sup} = \min_{E^a, E^b, D^a, D^b} \mathbb{E}_{x_{l,i}^a, x_{l,i}^b} [\mathcal{L}_{CE}(D^a(E^a(x_{l,i}^a)), y_i^a) + \mathcal{L}_{Dice}(D^b(E^b(x_{l,i}^b)), y_i^b)], \quad (2)$$

where $\mathcal{L}_{CE}$ and $\mathcal{L}_{Dice}$ represent cross-entropy loss function and Dice coefficient loss function. To further regularize underlying anatomical structures from a perspective of anatomical-wise predictions alignment on unlabeled data across modalities for multimodal segmentation, we introduce a Contrastive Anatomical-similar Consistency (CAC) loss. This loss is designed to measure the anatomical

similarity between predictions of different modalities on unlabeled data, thereby improving the accuracy and robustness of multimodal segmentation. The CAC loss incorporates the Dice coefficient similarity Sim_{Dice} [13] as a metric for assessing anatomical similarity. The $\mathcal{L}_{CAC}$ loss is defined as follows:

$$\mathcal{L}_{CAC} = \min_{E^a, E^b, D^a, D^b} \sum_j -\log \frac{\exp(Sim_{Dice}(D^a(E^a(x_{u,j}^a)), D^b(E^b(x_{u,j}^b))))}{\sum_{j=1}^M \exp(Sim_{Dice}(D^a(E^a(x_{u,j}^a)), D^b(E^b(x_{u,j}^b))))}. \quad (3)$$

Overall, the total loss of our framework is $\mathcal{L}_{total} = \mathcal{L}_{Sup} + \alpha\mathcal{L}_{CSC} + \beta\mathcal{L}_{CAC}$. To balance between these losses, we adopt a ramp-up weighting coefficient $\alpha = 0.1 \cdot e^{-5(1-\frac{t}{t_{max}})}$ and a ramp-down weighting coefficient $\beta = 0.1 \cdot e^{-5(\frac{t}{t_{max}})}$. This strategy follows related works [12, 18], where t and t_{max} denote the current and maximum number of epochs, respectively. All approaches were implemented using PyTorch on one NVIDIA A40 GPU. We trained the models using the Adam optimizer with an initial learning rate of 10^{-5} . The input size is $128 \times 128 \times 128$ voxels and batch size is 4.

3 Experiments

3.1 Datasets

We extensively evaluated our framework on three datasets of multi-modal image segmentation, including two public challenge datasets and one private dataset.

MS-CMRSeg Dataset. The MS-CMRSeg [11] dataset consists of cardiac magnetic resonance images from 45 patients with cardiomyopathy. The dataset includes three modalities: late gadolinium enhancement (LGE), T2-weighted (T2w) and balanced-steady state free precession (BSSFP). Notably, the dataset are paired as they originate from the same patients and are registered. The ground truth of the left ventricular cavity (LV), right ventricular cavity (RV), and left ventricular myocardium (Myo) is provided. For our experiment, we used 40 LGE/BSSFP image pairs for training and 5 for testing.

AMOS Dataset. The AMOS dataset [5] comes from a multi-modality abdominal multi-organ segmentation challenge. It includes 300 CT and 60 MRI images from multi-center, multi-modality and multi-phase, each annotations includes 15 abdominal organs. Notably, the CT/MRI data of AMOS dataset is unpaired as they originate from different patients. The training set consists of 200 CT and 40 MRI images, and the testing set comprises 100 CT and 20 MRI images. In our experiments, we randomly selected 40 CT images from the 200 available in the training set and included 40 MRI images, creating a new training set. Additionally, we randomly selected 20 CT images from the 100 available in the testing set and included 20 MRI images, forming a new testing set.

TAO Dataset. The thyroid-associated orbitopathy (TAO) dataset is an in-house multimodal dataset of thyroid-associated orbitopathy collected from the Gerald Choa Neuroscience Centre MRI Core Facility at the Prince of Wales Hospital in Hong Kong. This dataset comprises 100 cases, each of which underwent

orbital MRI and received a definitive TAO diagnosis. Each case includes pre-contrast T1-weighted (T1) and fat-suppressed post-contrast T1-weighted (T1c), which are sequentially acquired from the same patient. Notably, the T1/T1c data in the TAO dataset is unpaired due to differences in in-plane resolution. All MRI was performed on a 3.0 T Siemens scanner with a Head/Neck 64 Channel coil. T1 data was acquired using volumetric interpolated breath-hold examination (VIBE) pulse sequence with in-plane resolution of $0.555 \times 0.555 \text{ mm}^2$ and slice thickness of 1.5 mm. T1c data was collected using a fat-suppressed spoiled-gradient echo (GRE) core sequence with in-plane resolution of $0.625 \times 0.625 \text{ mm}^2$ and slice thickness of 1.5 mm. The manual annotation was performed by a trained rater using ITK-SNAP, under the guidance of a senior radiologist who has over 20 years of experience. The segmentation mask includes 8 anatomical structures of extraocular muscles, i.e., herniation of the lacrimal gland (LG), and compression and edema of the optic nerve (ON), inferior oblique muscle (IOM), superior oblique muscle (SOM), superior rectus (SR), lateral rectus (LR), medial rectus (MR), and inferior rectus (IR). For our experiments, we randomly selected 80 T1/T1c image pairs from 100 cases as the training set, and the remaining 20 T1/T1c image pairs as the testing set.

Table 1. Quantitative results of our approach and other methods on the AMOS [5] and TAO datasets in terms of Dice score.

AMOS 2022 dataset								
Methods	CT				MRI			
	Spleen	R. Kidney	L. Kidney	Liver	Spleen	R. Kidney	L. Kidney	Liver
Sup[7]	61.2±0.6	64.1±1.2	66.9±1.1	65.7±1.5	68.1±1.2	76.2±0.8	67.8±1.2	71.5±1.2
EFCD[4]	69.4±0.8	65.3±1.1	69.7±0.4	73.5±1.3	71.2±0.4	78.2±1.1	70.3±1.4	74.7±0.7
UMML[18]	71.4±1.2	66.3±1.2	70.2±1.3	75.3±0.7	75.8±1.7	76.7±0.8	72.1±1.9	76.2±1.5
mmFor[14]	72.6±1.3	73.5±0.5	73.1±1.4	79.7±1.7	76.3±1.4	78.9±1.2	72.2±1.5	76.4±1.8
CML[13]	73.6±0.8	75.4±1.3	76.2±1.7	81.2±1.1	77.8±1.6	81.1±0.7	73.9±1.7	77.8±1.4
Ours	74.8±0.9	76.5±1.3	77.9±1.8	80.7±1.8	78.4±1.8	79.5±1.5	74.7±1.8	79.6±1.3
TAO dataset								
Modality	Methods	SR	LG	IOM	LR	IR	MR	SOM
T1	Sup[7]	71.5±2.6	73.8±1.3	70.6±3.6	71.3±3.5	72.2±3.2	66.1±5.6	45.9±1.8
	EFCD[4]	73.1±8.8	75.6±1.8	71.2±0.7	72.8±1.5	74.3±2.7	68.8±4.7	51.7±9.1
	UMML[18]	73.4±3.2	77.3±3.2	73.7±1.3	74.5±1.1	75.2±4.2	69.1±0.4	52.3±1.4
	mmFor[14]	75.6±4.3	78.7±2.3	75.3±2.2	75.2±2.4	77.2±5.7	69.8±4.2	57.4±4.3
	CML[13]	76.7±3.1	80.6±5.4	79.7±3.3	78.8±2.3	78.3±3.3	71.5±6.2	62.4±5.8
	Ours	77.8±2.3	81.8±1.7	78.4±2.3	80.2±3.1	81.2±5.1	73.4±1.8	65.6±2.1
T1c	Sup[7]	65.5±3.6	67.3±4.2	66.6±2.2	68.3±2.5	69.7±3.2	62.1±4.4	47.5±4.5
	EFCD[4]	68.1±8.8	70.2±3.3	70.2±3.2	71.2±1.5	72.3±2.7	64.8±7.7	52.7±7.1
	UMML[18]	70.4±2.4	72.3±2.4	72.7±0.5	73.5±1.1	74.2±4.2	67.5±2.4	55.7±5.4
	mmFor[14]	73.6±3.5	73.7±0.5	73.3±0.2	75.3±0.4	76.2±5.7	70.4±4.2	59.4±5.3
	CML[13]	74.7±1.3	75.2±5.4	75.7±6.7	77.8±2.3	79.3±3.8	71.2±6.2	64.1±0.8
	Ours	75.9±4.3	77.8±3.3	77.4±8.1	79.2±5.3	79.1±5.1	72.4±4.7	66.3±5.4

3.2 Comparison with State-of-the-art Methods

We conducted extensive experiments, using 5% and 10% labeled data ratios from three datasets for training, and employed the Dice score and Average Symmetric Surface Distance (ASSD) for quantitative evaluation. We compared our model




Fig. 2. Comparison of segmentation performance across different models on the MS-CMRSeg [11] dataset using various labeled data ratios, evaluated in terms of Dice score.

with other multimodal learning methods, such as EFCD [4] and mmFor [14]. Additionally, we compared our model with semi-supervised multimodal learning approaches, such as UMML [18] and CML [13], and the fully supervised method V-Net [7]. Table 1 presents the quantitative performance of different methods on the AMOS [5] and TAO datasets. The results show that our framework considerably surpasses the comparison methods in both CT and MRI modalities, achieving high Dice scores in a label-scarce scenario on the AMOS [5]. As is well known, unpaired data from different modalities originate from different patients and cannot be directly aligned. This complicates the extraction of consistent features from different modalities compared to paired data. Despite the AMOS dataset [5] is unpaired, our model demonstrates superior performance. This can be attributed to its architecture, which is based on the Transformer model of SAM-Med3D [10]. The input images are divided into patches, which are then linearly embedded and combined with positional encoding, providing strong feature learning capability. We then utilize the CSC loss to align the channel-wise features from multimodal images. Furthermore, we introduce a novel MIA module to effectively harness modality-independent knowledge from each modality, facilitating efficient feature fusion. Consequently, our model achieves superior performance in semi-supervised multimodal segmentation tasks. For the TAO dataset, specifically for the T1 and T1c modalities, our framework has also shown promising results. Fig. 2 provides a comparative analysis of segmentation performance across various models on the MS-CMRSeg dataset. The performance results evaluated using the ASSD score on the MS-CMRSeg dataset, are presented in Table 3 of the supplementary material. It is evident from the results

that our method surpasses the comparison methods in both BSSFP and LGE modalities, consistently achieving high Dice scores.




Fig. 3. (a) and (b) Visual comparison between the CML [13] and our method on the TAO and MS-CMRSeg [11] datasets. (c) Ablation analysis of the cross CMC strategy. (d) Ablation study of single and multiple modalities for training our model on the MS-CMRSeg dataset [11].

Fig. 3 (a) and (b) present a qualitative comparison of our framework with the CML [13] method for thyroid-associated orbitopathy and cardiac segmentation. The results demonstrate that our model’s predictions align more closely with the ground truth and contain fewer erroneous regions compared to CML [13].

Table 2. Ablation results of our method on the AMOS dataset [5].

Methods				Dice (%)				ASSD (mm)			
				CT		MRI		CT		MRI	
Baseline	Modality-Specific Encoder	CMC Strategy	CCL Module	Spleen	Liver	Spleen	Liver	Spleen	Liver	Spleen	Liver
✓				70.4	76.7	73.3	74.4	4.41	4.12	4.38	4.16
✓	✓			72.9	78.5	75.9	76.2	4.17	3.85	3.87	3.73
✓	✓	✓		73.4	79.2	76.2	77.3	3.58	3.47	3.49	3.38
✓	✓	✓	✓	74.8	80.7	78.4	79.6	3.42	3.38	3.23	3.17

3.3 Ablation Study

We investigate the effectiveness of three key components in our method: modality-specific encoder, cross modality collaboration (CMC) strategy, and contrastive consistent learning (CCL) module. As shown in Table 2, we first set up a baseline network without using these components. Subsequently, we add the modality-specific encoder, CMC strategy, and CCL module one by one into the baseline network. Integrating these components consistently enhances performance, thereby highlighting their critical significance within our framework. Fig. 3 (c) presents an ablation analysis of our framework on the TAO dataset, comparing the performance with and without using the CMC strategy. Additionally,

we conduct an ablation study on the MS-CMRSeg [11] dataset to compare the performance of using single and multiple modalities for training our model, as shown in Fig. 3 (d). The results demonstrate that using multiple modalities (i.e., BSSF+LGE) outperforms using a single modality (i.e., BSSF or LGE). The integration of the CMC strategy into our framework significantly enhances the highlighting of object regions and improves the capacity for class-discriminative localization compared to frameworks without the CMC strategy. Consequently, this demonstrates that the CMC strategy is able to foster a more efficient feature fusion.

4 Conclusion

We propose a novel semi-supervised multimodal segmentation framework that synergistically incorporates cross-modality collaboration and contrastive consistent learning strategies to achieve modality-independent knowledge and feature alignment across different modalities. We validate our method on three benchmark datasets and achieve comparable results to the state-of-the-art approaches. Notably, the framework exhibits exceptional robustness against significant inference variations, underscoring its potential for widespread clinical application in real-world clinical settings. More efficient fusion strategies on 3D medical image datasets and cross-modality settings will be conducted in our future studies.

Acknowledgements. This work was supported by Hong Kong Research Grants Council under Projects No. T45-401/22-N and Project No. 14200721.

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

References

1. Basak, H., Yin, Z.: Pseudo-label guided contrastive learning for semi-supervised medical image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19786–19797 (2023)
2. Chen, X., Zhou, H.Y., Liu, F., Guo, J., Wang, L., Yu, Y.: Mass: modality-collaborative semi-supervised segmentation by exploiting cross-modal consistency from unpaired ct and mri images. *Med. Image Anal.* **80**, 102506 (2022)
3. Huang, L., Denoeux, T., Vera, P., Ruan, S.: Evidence fusion with contextual discounting for multi-modality medical image segmentation. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) International Conference on Medical Image Computing and Computer-Assisted Intervention. LNCS. vol. 13435, pp. 401–411. Springer, Cham (2022)
4. Huang, L., Denoeux, T., Vera, P., Ruan, S.: Evidence fusion with contextual discounting for multi-modality medical image segmentation. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) International Conference on Medical Image Computing and Computer-Assisted Intervention. LNCS. vol. 13435, pp. 401–411. Springer, Cham (2022)

5. Ji, Y., Bai, H., Ge, C., Yang, J., Zhu, Y., Zhang, R., Li, Z., Zhanng, L., Ma, W., Wan, X., et al.: Amos: a large-scale abdominal multi-organ benchmark for versatile medical image segmentation. *Advances in Neural Information Processing Systems* **35**, 36722–36732 (2022)
6. Liu, J., Zhang, Y., Chen, J.N., Xiao, J., Lu, Y., A Landman, B., Yuan, Y., Yuille, A., Tang, Y., Zhou, Z.: Clip-driven universal model for organ segmentation and tumor detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 21152–21164 (2023)
7. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: fully convolutional neural networks for volumetric medical image segmentation. In: *2016 Fourth International Conference on 3D Vision (3DV)*. pp. 565–571. IEEE (2016)
8. Qiu, J., Li, L., Wang, S., Zhang, K., Chen, Y., Yang, S., Zhuang, X.: Myops-net: myocardial pathology segmentation with flexible combination of multi-sequence cmr images. *Med. Image Anal.* **84**, 102694 (2023)
9. Rahutomo, F., Kitasuka, T., Aritsugi, M., et al.: Semantic cosine similarity. In: *The 7th international student conference on advanced science and technology ICAST*. vol. 4, p. 1. University of Seoul South Korea (2012)
10. Wang, H., Guo, S., Ye, J., Deng, Z., Cheng, J., Li, T., Chen, J., Su, Y., Huang, Z., Shen, Y., et al.: Sam-med3d. *arXiv preprint arXiv:2310.15161* (2023)
11. Wu, F., Zhuang, X.: Minimizing estimated risks on unlabeled data: A new formulation for semi-supervised medical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(5), 6021–6036 (2022)
12. Yu, L., Wang, S., Li, X., Fu, C.W., Heng, P.A.: Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In: Shen, D., et al. (eds.) *International Conference on Medical Image Computing and Computer-Assisted Intervention*. LNCS. vol. 11766, pp. 605–613. Springer, Cham (2019)
13. Zhang, S., Zhang, J., Tian, B., Lukasiewicz, T., Xu, Z.: Multi-modal contrastive mutual learning and pseudo-label re-learning for semi-supervised medical image segmentation. *Med. Image Anal.* **83**, 102656 (2023)
14. Zhang, Y., He, N., Yang, J., Li, Y., Wei, D., Huang, Y., Zhang, Y., He, Z., Zheng, Y.: mmformer: Multimodal medical transformer for incomplete multimodal learning of brain tumor segmentation. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) *International Conference on Medical Image Computing and Computer-Assisted Intervention*. LNCS. vol. 13435, pp. 107–117. Springer, Cham (2022)
15. Zhang, Y., Yang, J., Tian, J., Shi, Z., Zhong, C., Zhang, Y., He, Z.: Modality-aware mutual learning for multi-modal medical image segmentation. In: de Bruijne, M., et al. (eds.) *International Conference on Medical Image Computing and Computer-Assisted Intervention*. LNCS. vol. 12902, pp. 589–599. Springer, Cham (2021)
16. Zhou, X., Li, Z., Tong, T.: Dtsc-net: semi-supervised 3d biomedical image segmentation through dual-teacher simplified consistency. In: *2022 IEEE international conference on bioinformatics and biomedicine*. pp. 1429–1434. IEEE (2022)
17. Zhou, X., Li, Z., Xue, Y., Chen, S., Zheng, M., Chen, C., Yu, Y., Nie, X., Lin, X., Wang, L., et al.: Cuss-net: A cascaded unsupervised-based strategy and supervised network for biomedical image diagnosis and segmentation. *IEEE J. Biomed. Health Inform* **27**(5), 2444 – 2455 (2023)
18. Zhu, L., Yang, K., Zhang, M., Chan, L.L., Ng, T.K., Ooi, B.C.: Semi-supervised unpaired multi-modal learning for label-efficient medical image segmentation. In: de Bruijne, M., et al. (eds.) *International Conference on Medical Image Computing and Computer-Assisted Intervention*. LNCS. vol. 12902, pp. 394–404. Springer, Cham (2021)