

Hypergraph Tversky-Aware Domain Incremental Learning for Brain Tumor Segmentation with Missing Modalities

Junze Wang^{1†}, Lei Fan^{2,3†}, Weipeng Jing¹, Donglin Di⁴, Yang Song³, Sidong Liu⁵, and Cong Cong⁵✉

¹ College of Computer and Control Engineering, Northeast Forestry University, Harbin, China

² The Centre for Healthy Brain Ageing (CHeBA), UNSW, Sydney, Australia

³ School of Computer Science and Engineering, UNSW, Sydney, Australia

⁴ School of Software, Tsinghua University, Beijing, China

⁵ Centre for Health Informatics, Macquarie University, Sydney, Australia
thomas.cong@mq.edu.au

Abstract. Existing methods for multimodal MRI segmentation with missing modalities typically assume that all MRI modalities are available during training. However, in clinical practice, some modalities may be missing due to the sequential nature of MRI acquisition, leading to performance degradation. Furthermore, retraining models to accommodate newly available modalities can be inefficient and may cause overfitting, potentially compromising previously learned knowledge. To address these challenges, we propose Replay-based Hypergraph Domain Incremental Learning (ReHyDIL) for brain tumor segmentation with missing modalities. ReHyDIL leverages Domain Incremental Learning (DIL) to enable the segmentation model to learn from newly acquired MRI modalities without forgetting previously learned information. To enhance segmentation performance across diverse patient scenarios, we introduce the Cross-Patient Hypergraph Segmentation Network (CHSNet), which utilizes hypergraphs to capture high-order associations between patients. Additionally, we incorporate Tversky-Aware Contrastive (TAC) loss to effectively mitigate information imbalance both across and within different modalities. Extensive experiments on the BraTS2019 dataset demonstrate that ReHyDIL outperforms state-of-the-art methods, achieving an improvement of over 2% in the Dice Similarity Coefficient across various tumor regions. Our code is available at ReHyDIL.

Keywords: Missing modality · Brain tumor segmentation · Domain incremental learning · Class imbalance · Hypergraph.

[†] Equal contribution

✉ Corresponding author

1 Introduction

In clinical practice, brain tumor diagnosis is typically supported by observations using four routine Magnetic Resonance Imaging (MRI) modalities: T1, T1CE, T2, and FLAIR [13]. Each modality provides unique tissue contrasts that aid in identifying different tumor regions, such as the whole tumor (WT), tumor core (TC), and enhancing tumor (ET) [16]. Consequently, accurate and automated segmentation of these modalities is crucial for effective diagnosis and treatment.

Recent advances in deep learning have significantly enhanced multimodal brain tumor segmentation [10,19]. However, model performance often degrades due to missing modalities, which is a common issue caused by clinical constraints such as varying acquisition protocols [30]. Existing methods address this issue primarily through feature fusion [26,29,5,28], knowledge distillation [12,17,27], and modality masking [22,25,18]. However, these methods often assume that all MRI modalities are available during training, which is a condition often unmet in practice. To address this challenge, we draw inspiration from Domain Incremental Learning (DIL) [21,24]. However, directly applying DIL presents challenges. Specifically, catastrophic forgetting occurs when newly added modalities overshadow previously acquired ones during training, leading to an inter-modality imbalance. Moreover, tumor regions typically occupy only a small portion of the image relative to the background, resulting in an intra-modality imbalance [3,4]. These issues hinder the model’s ability to learn robust cross-modal features, ultimately degrading segmentation performance in missing modality scenarios.

Accordingly, we introduce the Replay-based Hypergraph Domain Incremental Learning (ReHyDIL). ReHyDIL employs a replay buffer to retain data from previous modalities, enabling the model to continue domain adaptation [7,8], learn cross-modal features and mitigate catastrophic forgetting. To address inter-modality imbalance, we propose the Tversky-Aware Contrastive (TAC) loss with a balanced queue. This loss encourages the model to learn robust correlations between different modalities, bridging the information gap caused by missing modalities during inference and reducing modality imbalance. For intra-modality imbalance at the pixel level, we employ Tversky loss and Tversky-based Focal loss, ensuring better segmentation of small and less-represented tumor structures. Moreover, to effectively leverage the rich patient-level information in MRI datasets, we propose CHSNet, a hypergraph-based segmentation network designed to capture complex inter-patient relationships. In CHSNet, each patient is represented as a vertex, connected by hyperedges to others with similar tumor features. The contributions can be summarized as follows:

- We propose ReHyDIL for brain tumor segmentation with missing modality in scenarios where MRI modalities are incrementally available during training.
- We introduce the Tversky-Aware Contrastive (TAC) loss to balance the learning between different modalities.
- We present the Cross-Patient Hypergraph Segmentation Network (CHSNet), which captures high-order associations across patients, enhancing feature representations and improving segmentation performance.




Fig. 1: The ReHyDIL pipeline processes a set of MRI modalities acquired sequentially, leveraging incremental learning to train CHSNet. This training is guided by our proposed TAC loss, alongside other Tversky-based segmentation losses. At the end of each learning phase, a replay buffer is constructed to store samples from the currently learned modalities. This mechanism enables the trained CHSNet to effectively capture cross-modality features, improving the segmentation of different tumor regions even in the presence of missing modalities.

- Compared to existing state-of-the-art methods, ReHyDIL achieves higher Dice Similarity Coefficients (DSC) on the BraTS2019 dataset, with gains of 1.05% for Whole Tumor (WT), 2.53% for Tumor Core (TC), and 1.52% for Enhancing Tumor (ET).

2 Method

2.1 Replay-based Hypergraph Domain Incremental Learning

We illustrate the overall workflow of ReHyDIL in Fig. 1, where the core objective is to train the proposed CHSNet (Section 2.2) $F_{\theta^m} : X_p^m \rightarrow Y_p^{c,m}$ with m continuously arriving MRI modalities. In this study, each consecutively arriving MRI modality strictly corresponds to one training stage. X_p^m and $Y_p^{c,m}$ denote the input sample and its prediction for patient p , and c denotes the tumor region label. Initially, F_{θ^1} is trained using samples from the first modality D_1 , and a replay buffer is initialized. In the subsequent training stages ($1 < i \leq m$), where new modalities are introduced, F_{θ^i} loads the weights from the previous stage and is trained using the proposed TAC loss (Section 2.3) along with samples from previously encountered modalities stored in the replay buffer.

To build the replay buffer R with a capacity T , we retain $\mathcal{P}\%$ of samples from each training stage D_i . Then T is calculated as $T = \sum_{i=1}^m \|D_i\| \times \mathcal{P}\%$. For a given modality D_i , we follow the mean loss strategy [11] to select the top $\mathcal{P}\%$ of samples with loss values closest to the median loss μ . Specifically, for




Fig. 2: Illustration of the proposed CHSNet. CHSNet follows a U-Net architecture, incorporating hypergraphs into the last and second-to-last encoder layers. These hypergraphs effectively encode cross-patient information into the output features, which serve as inputs to the subsequent decoder blocks.

each sample X_p^i , we compute its deviation: $d_p = |\mathcal{L}_{total}(X_p^i, \theta^i) - \mu|$. Samples are ranked in ascending order based on d_p , with $\sigma(X_p^i)$ denoting the rank of X_p^i . The top $\mathcal{P}\%$ ranked samples are stored in the replay buffer according to:

$$R = \{X_p^i \in D_i \mid \sigma(X_p^i) \leq \|D_i\| \times \mathcal{P}\%\}. \quad (1)$$

Samples stored in R will be used together with D_i to train CHSNet using TAC loss. This enables the final model, F_{θ^m} , to incrementally learn from all modalities and effectively perform segmentation even in the presence of missing modalities.

2.2 Cross-Patient Hypergraph Segmentation Network

CHSNet is designed to capture cross-patient associations and enhance segmentation performance by integrating a Cross-Patient Hypergraph (CPH). To enforce this cross-patient learning, we impose a constraint during each batch sampling to ensure that selected samples originate from different patients. As illustrated in Fig. 2, CHSNet is built upon the U-Net architecture [14]. Empirically, we embed the hypergraph into the 4th and 5th encoders to learn different feature representations, enabling more robust segmentation performance.

We construct each hypergraph using the encoder’s output, $\mathbf{E}_f \in \mathbb{R}^{B \times C \times H \times W}$. First, we flatten $\mathbf{E}_f$ from (B, C, H, W) to (N, C) , where $N = B \times H \times W$ represents the total number of hypergraph vertices. Each vertex corresponds to a C -dimensional feature vector, formally denoted as: $\mathbf{V} = [\mathbf{v}_1; \mathbf{v}_2; \dots; \mathbf{v}_N] \in \mathbb{R}^{N \times C}$, where $\mathbf{v}_j \in \mathbb{R}^C$. To construct the hypergraph $\mathbf{G}$, we employ a **K**-Nearest Neighbours (**KNN**) method based on Euclidean distance [9,23]. Initially, a **KNN** search is performed for each vertex with $\mathbf{K} = 1$, generating a preliminary incidence matrix $\mathbf{H}$, where a value of 1 indicates a connection between vertices via a hyperedge. Next, we refine the hypergraph structure by updating the $\mathbf{K}$ -value of each vertex $\mathbf{v}_j$ to match its degree in $\mathbf{H}$, i.e., $\mathbf{K}(j)$ is set to the degree of $\mathbf{v}_j$. A second **KNN** search is then performed using the updated $\mathbf{K}$ -values to establish the final hypergraph structure. As a result, each vertex $\mathbf{v}_j$ is connected to its set of nearest neighbors, $\mathbf{N}_{\mathbf{K}(j)}$, defining the hyperedge

as: $\mathbf{e}_j = \{\mathbf{v}_j\} \cup \mathbf{N}_{\mathbf{K}(j)}(\mathbf{v}_j)$. This process yields the updated incidence matrix $\mathbf{H} \in \mathbb{R}^{N \times K}$, along with the vertex degree matrix $\mathbf{D}_v$ and hyperedge degree matrix $\mathbf{D}_e$.

Subsequently, HGNN [12] is employed to integrate cross-patient higher-order relationships. The updated hypergraph features $\mathbf{V}'$ are computed as follows:

$$\mathbf{V}' = \mathbf{D}_v^{-\frac{1}{2}} \mathbf{H} \mathbf{W}_e \mathbf{D}_e^{-1} \mathbf{H}^\top \mathbf{D}_v^{-\frac{1}{2}} \mathbf{V}, \quad (2)$$

where $\mathbf{D}_v^{-\frac{1}{2}}$ normalizes vertex degrees, $\mathbf{W}_e$ is a learnable weight for hyperedges and is initially set to 1. Next, we reshape $\mathbf{V}'$ back to (B, C, H, W) , aligning it with the original feature map $\mathbf{E}_f$. Finally, we concatenate $\mathbf{V}'$ and $\mathbf{E}_f$ along the channel dimension and fuse the resulting features via a 1×1 convolution, producing integrated feature representations to be passed into subsequent decoder blocks.

2.3 Tversky-Aware Contrastive Loss

We propose the TAC loss to regularize CHSNet by leveraging samples stored in R alongside those from the current modality D_i . However, due to the limited size of R , which can be easily overwhelmed by D_i , we introduce a balanced queue (Q) mechanism to effectively address this inter-modality imbalance. Q is designed as a double-queue structure, where one queue (Q_R) stores predictions of samples from R generated by $F_{\theta^{i-1}}$, and the other (Q_D) contains predictions of samples from D_i produced by F_{θ^i} .

Specifically, given a batch of B predictions, we randomly select S predictions, where $S = \min(B, \|R\|)$, from each Q_R and Q_D , and normalize them before storing them in the memory bank Q , ensuring that $\|Q\| = 2S$. For each prediction $Y_p^c \in Q_i$, we define its positive pair $Y_q^c \in Q_{j \neq i}$. This means that they come from different patients and different modalities but share the same predicted region. Similarly, we define the negative pair $Y_q^z \in Q_{j \neq i}$, where $p \neq q$ and $c \neq z$. Then contrastive learning is conducted via:

$$\mathcal{L}_{\text{TAC}} = -\log \frac{\exp(S_{\text{tve}}(Y_p^c, Y_q^c)/\tau)}{\exp(S_{\text{tve}}(Y_p^c, Y_q^c)/\tau) + \sum_{p \neq q, c \neq z} \exp(S_{\text{tve}}(Y_p^c, Y_q^z)/\tau)}, \quad (3)$$

where τ is the temperature coefficient and is set to 1, S_{tve} is the Tversky similarity. Following [20], we define S_{tve} as:

$$S_{\text{tve}}(g, u) = \frac{gu}{gu + \alpha g(1 - u) + \beta(1 - g)u}, \quad (4)$$

where g and u denote two different predictions, α and β are hyperparameters that control the trade-off between false positives (FPs, i.e., false detections) and false negatives (FNs, i.e., missed detections). In Equation 3, the anchor prediction Y_p^c is used as g , while the predictions of the paired samples Y_q^c, Y_q^z are used as u .

However, due to intra-modality imbalance, a naively trained CHSNet may produce high-precision but low-recall segmentations, where $\text{FNs} \gg \text{FPs}$. To mitigate this imbalance, we explicitly set $\beta > \alpha$, applying a greater penalty to

Table 1: Segmentation results of WT, TC, and ET in terms of DSC(%) on the BraTS2019 dataset. Missing and available modalities are denoted by $\circ$ and $\bullet$, respectively. “ $\dagger$ ” denotes significance level is reached as p-value < 0.05 .

Modalities				WT					TC					ET				
T2	T1c	T1	FL	mmF	MFI	M ³	M+P	Ours	mmF	MFI	M ³	M+P	Ours	mmF	MFI	M ³	M+P	Ours
$\bullet$	$\circ$	$\circ$	$\circ$	80.66	81.87	84.49	82.74	86.71$\dagger$	69.45	58.26	71.35	55.74	73.68$\dagger$	37.93	39.82	47.61	36.29	49.75$\dagger$
$\circ$	$\bullet$	$\circ$	$\circ$	73.64	74.39	78.62	75.88	79.48$\dagger$	78.62	81.32	83.95	82.83	84.65$\dagger$	70.64	76.53	79.33	79.29	80.17$\dagger$
$\circ$	$\circ$	$\bullet$	$\circ$	73.76	73.28	79.32	78.32	80.77$\dagger$	61.36	65.14	65.98	62.13	69.34$\dagger$	33.05	33.82	44.43	38.40	45.22$\dagger$
$\circ$	$\circ$	$\circ$	$\bullet$	87.31	86.96	86.38	87.05	88.25$\dagger$	66.85	57.31	70.37	64.53	72.87$\dagger$	37.37	40.67	45.62	38.23	48.61$\dagger$
$\bullet$	$\bullet$	$\circ$	$\circ$	84.55	85.81	85.74	84.47	87.13$\dagger$	79.98	85.20	83.25	82.94	85.96$\dagger$	71.92	80.04	78.76	79.75	80.51$\dagger$
$\circ$	$\bullet$	$\bullet$	$\circ$	77.60	74.21	80.31	79.95	81.48$\dagger$	79.98	83.09	85.43	86.59	86.32	71.38	77.47	80.31	82.35	80.97
$\circ$	$\circ$	$\bullet$	$\bullet$	88.90	88.09	90.17	89.12	89.63$\dagger$	70.44	65.96	66.32	66.59	71.75$\dagger$	41.93	41.47	48.83	42.49	49.55$\dagger$
$\bullet$	$\circ$	$\bullet$	$\circ$	83.17	84.52	86.29	85.30	87.58$\dagger$	68.73	60.63	65.52	62.16	69.42$\dagger$	42.46	43.36	44.65	38.98	47.36$\dagger$
$\bullet$	$\circ$	$\circ$	$\bullet$	88.68	89.53	90.21	88.67	90.39	70.72	65.43	70.21	68.48	70.96$\dagger$	43.54	47.29	46.69	46.75	46.93
$\circ$	$\bullet$	$\circ$	$\bullet$	88.50	89.91	90.75	89.56	91.42$\dagger$	80.84	82.26	83.02	82.37	83.22$\dagger$	72.03	77.52	78.17	75.58	78.65$\dagger$
$\circ$	$\bullet$	$\bullet$	$\bullet$	89.02	89.78	89.89	89.91	90.75$\dagger$	82.33	85.96	86.13	86.06	86.46	72.82	83.23	84.17	83.19	83.26
$\bullet$	$\circ$	$\bullet$	$\bullet$	89.08	90.06	90.44	89.99	91.19$\dagger$	71.90	66.37	71.87	59.87	72.50$\dagger$	45.34	47.89	50.26	37.08	48.11$\dagger$
$\bullet$	$\bullet$	$\circ$	$\bullet$	88.94	91.05	91.21	89.83	91.46$\dagger$	81.64	86.99	86.53	86.20	87.79$\dagger$	72.25	83.32	82.48	82.10	84.62$\dagger$
$\bullet$	$\bullet$	$\bullet$	$\circ$	84.95	86.47	87.89	86.59	87.98$\dagger$	81.27	85.58	86.41	88.12	89.24$\dagger$	72.98	81.11	82.35	82.88	83.46$\dagger$
$\bullet$	$\bullet$	$\bullet$	$\bullet$	89.21	90.81	91.36	90.35	92.43$\dagger$	82.47	86.79	86.51	87.73	88.06$\dagger$	72.74	83.34	83.88	82.76	85.35$\dagger$
Means				84.53	85.12	86.87	85.91	87.78	75.09	74.42	77.53	74.82	79.49	57.23	62.46	65.17	61.74	66.16

FNs to better capture rare positive objects, as shown in Fig. 1. Furthermore, we design an intra-modality regularization $\mathcal{L}_{\text{intra}} = \mathcal{L}_{\text{DT}} + \mathcal{L}_{\text{FT}}$ under this setup ($\beta > \alpha$). Specifically, $\mathcal{L}_{\text{DT}}$ is a Tversky-based Dice loss, which extends Dice loss with Tversky similarity to systematically suppress FNs. Furthermore, $\mathcal{L}_{\text{FT}}$ integrates Tversky similarity into Focal loss [15] to further reduce FNs by focusing on difficult-to-classify samples:

$$\mathcal{L}_{\text{DT}}(y, t) = 1 - \sum_{i=1}^n S_{\text{tve}}(y_i, t_i), \quad \mathcal{L}_{\text{FT}}(y, t) = (1 - \sum_{i=1}^n S_{\text{tve}}(y_i, t_i))^\gamma, \quad (5)$$

where y_i denotes the i -th predicted pixel probability, t denotes its ground truth label, and γ is used to adjust the level of attention given to minority class pixels. The total loss is supervised by $\mathcal{L}_{\text{TAC}}$, and $\mathcal{L}_{\text{intra}}$:

$$\mathcal{L}_{\text{total}} = \omega \mathcal{L}_{\text{TAC}} + \mathcal{L}_{\text{intra}}. \quad (6)$$

3 Experiment

Datasets. BraTS2019 dataset [1] is a multimodal MRI dataset specifically designed for brain tumor segmentation, comprising 335 patient-level 3D image volumes. Each case includes four modalities: T1, T1CE (T1c), T2, and FLAIR (FL), and three tumor regions: ET, TC, and WT. For each 3D volume, 2D slices were extracted along the z-axis and cropped to a resolution of 224×224 pixels. The dataset was divided into training, validation, and testing subsets using an 8:1:1 ratio, resulting in 41,540 training slices from 268 patients, 5,115 validation




Fig. 3: The segmentation visualizations of Dice-based method and ReHyDIL (Tversky similarity), where the difference is only the similarity used. WT, TC, and ET are marked with light gray, medium gray and white, respectively.

slices from 33 patients, and 5,270 testing slices from 34 patients. To maintain data integrity, slices from the same patient were assigned to the same subset.

Experimental Setup. All models were trained using PyTorch on $2 \times$ NVIDIA Tesla V100 GPUs for 200 epochs with a batch size of 64. The Adam optimizer was employed with a learning rate of 1×10^{-4} and a weight decay of 4×10^{-4} . A warm-up multi-step learning rate scheduler with four steps was utilised. For hyperparameters, we set $\omega = 0$ for D_1 ; for all subsequent modalities, ω was set to 1. We set $\mathcal{P} = 10$, $\alpha = 0.7$, $\beta = 1.5$, and $\gamma = 1.2$. Furthermore, the modalities were utilized sequentially for training, following the typical order of clinical acquisition [6]: $T1 \rightarrow T2 \rightarrow FLAIR \rightarrow T1CE$. Missing modality scenarios were implemented by setting D_i (where $i \in \{T2, T1CE, T1, FLAIR\}$) to 0. For the evaluation metric, we calculated the Dice Similarity Coefficient (DSC) for each tumor region within each patient’s 3D image volume and compared our method with the second-best method using a paired t-test to obtain p-values.

Quantitative Results. We compared the proposed method with state-of-the-art approaches for missing modality brain tumor segmentation, including feature-fusion methods mmFormer (mmF) [26] and UNet-MFI (MFI) [29], knowledge distillation methods M^3AE (M^3) [12] and PASSION [17], and the modality-masked method $M^2FTrans$ (M+P) [18]. Specifically, PASSION is applied to $M^2FTrans$ (M+P) as described in [17]. As shown in Table 1, ReHyDIL outperforms all compared methods in 41 out of 45 cases, demonstrating a superior average Dice Similarity Coefficient (DSC). Specifically, for Whole Tumor (WT), ReHyDIL achieves an average DSC of 87.78%, representing a 3.13% improvement over MFI [29]. For Tumor Core (TC), it attains an average DSC of 79.49%, reflecting a 2.53% gain relative to M^3 [12]. Regarding Enhancing Tumor (ET), ReHyDIL reaches an average DSC of 66.16%, signifying a 7.16% enhancement over M+P [18]. Moreover, the p-values are less than 0.05 in 39 out of 45 cases, indicating that the performance improvements of ReHyDIL over other baselines are statistically significant. It is noteworthy that all compared methods assume the availability of all MRI modalities during training. In contrast, our method

Table 2: Ablation study for (a) the encoder layers where CPH is added and different similarity measure in contrastive losses, (b) Tversky similarity hyperparameters α and β , (c) different modules in ReHyDIL.

		DSC(%)					DSC(%)					DSC(%)			
		WT	TC	ET	α	β	WT	TC	ET	DIL	CPH	TAC	WT	TC	ET
Layers	5	85.32	76.48	61.95	0.5	0.5	80.12	74.58	60.31				60.14	44.62	33.58
	5, 4	87.78	79.49	66.16	0.9	1.3	84.73	77.04	63.19	✓			76.42	55.04	47.29
	5, 4, 3	79.95	68.59	50.82	0.8	1.4	85.96	77.27	64.85	✓	✓		82.56	71.85	58.76
Similarity	Cosine	84.18	75.03	62.74	0.7	1.5	87.78	79.49	66.16	✓		✓	82.56	71.85	58.76
	Tversky	87.78	79.49	66.16	0.6	1.6	85.29	76.71	62.48	✓	✓	✓	87.78	79.49	66.16

(a)

(b)

(c)

(a)

(b)

(c)

operates under a more realistic scenario where modalities are obtained incrementally. The superior performance of ReHyDIL further demonstrates its potential as a reliable tool for real clinical applications.

Qualitative Results. The segmentation visualizations are presented in Fig. 3. Integrating additional modalities provides complementary features, which enhance segmentation performance. For example, when only the FLAIR modality is available, ReHyDIL can roughly delineate the whole tumor (WT) but struggles to accurately segment the enhancing tumor (ET) and tumor core (TC). Incorporating T1CE and T2 modalities (i.e., FLAIR+T1CE+T2) significantly improves ET segmentation, achieving a closer alignment with the ground truth (GT). Furthermore, the visualizations demonstrate that Tversky similarity outperforms Dice similarity in segmenting boundary regions and smaller regions, highlighting its effectiveness in mitigating intra-modality imbalance.

Ablation Study. We first compared the impact of adding Cross-Patient Hypergraph (CPH) modules to different layers of CHSNet, as shown in Table 2(a). The results indicate that two CPH layers yield the best segmentation performance, while adding a third layer causes a significant decline due to over-smoothing. Additionally, we examined the effectiveness of different similarity metrics in the TAC loss. Compared to cosine similarity [2], using Tversky similarity improves the Dice score, as it directly measures the overlap between target regions, effectively mitigating intra-modality imbalance.

Furthermore, the ablation studies reveal a key trade-off in controlling false negatives (FNs), as shown in Table 2(b). With $\alpha = \beta = 0.5$, Tversky similarity S_{tve} simplifies to the Dice similarity, which performs poorly due to an insufficient FN penalty. Gradually increasing β from 1.3 to 1.5 significantly improves DSC across all regions, highlighting the importance of penalizing FNs to reduce missed detections. The best performance is achieved at $\alpha = 0.7, \beta = 1.5$. However, further increasing β to 1.6 leads to performance collapse, likely from over-suppressing ambiguous boundaries. Finally, we conducted ablation studies on the overall framework. Without CPH, CHSNet reduces to the U-Net, and without TAC loss, the model relies solely on intra-modality regularization $\mathcal{L}_{\text{intra}}$. These results confirm that both CPH and TAC loss individually and collectively enhance segmentation performance—CPH by integrating high-order inter-patient relationships and TAC loss by bridging the information gap during inference.

4 Conclusion

In this paper, we propose the ReHyDIL framework for brain tumor segmentation with missing modalities, leveraging DIL to incrementally learn from new modalities while preserving prior knowledge. Specifically, we propose TAC loss, which effectively balances the learning process across different modalities and mitigates the information gap caused by missing modalities during inference. Additionally, we introduce CHSNet to capture high-order relationships between patients, enhancing segmentation performance. Extensive experiments on the BraTS2019 dataset demonstrate that ReHyDIL achieves state-of-the-art performance in missing modality scenarios. Though effective, ReHyDIL has certain limitations. Its use of a replay buffer introduces additional memory overhead, and the framework is currently limited to four MRI modalities with similar image characteristics. In future work, we aim to explore alternative DIL strategies and develop more advanced alignment mechanisms or adaptors to effectively fuse features from diverse modalities.

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

Acknowledgments. This work was partially supported by an NHMRC Ideas Grant (GNT202035).

References

1. Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R.T., Berger, C., Ha, S.M., Rozycki, M., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629* (2018)
2. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *ICML*. pp. 1597–1607. PmLR (2020)
3. Cong, C., Liu, S., Rana, P., Pagnucco, M., Di Ieva, A., Berkovsky, S., Song, Y.: Adaptive unified contrastive learning with graph-based feature aggregator for imbalanced medical image classification. *Expert Systems with Applications* **251**, 123783 (2024)
4. Cong, C., Xuan, S., Liu, S., Zhang, S., Pagnucco, M., Song, Y.: Decoupled optimisation for long-tailed visual recognition. In: *AAAI*. vol. 38, pp. 1380–1388 (2024)
5. Ding, Y., Yu, X., Yang, Y.: Rfnet: Region-aware fusion network for incomplete multi-modal brain tumor segmentation. In: *ICCV*. pp. 3975–3984 (2021)
6. Ellingson, B.M., Bendszus, M., Boxerman, J., Barboriak, D., Erickson, B.J., Smits, M., Nelson, S.J., Gerstner, E., Alexander, B., Goldmacher, G., et al.: Consensus recommendations for a standardized brain tumor imaging protocol in clinical trials. *Neuro-oncology* **17**(9), 1188–1198 (2015)
7. Fan, L., Ding, Y., Fan, D., Di, D., Pagnucco, M., Song, Y.: Grainspace: A large-scale dataset for fine-grained and domain-adaptive recognition of cereal grains. In: *CVPR*. pp. 21116–21125 (2022)

8. Fan, L., Sowmya, A., Meijering, E., Song, Y.: Fast ff-to-ffpe whole slide image translation via laplacian pyramid and contrastive learning. In: MICCAI. pp. 409–419. Springer (2022)
9. Gao, Y., Zhang, Z., Lin, H., Zhao, X., Du, S., Zou, C.: Hypergraph learning: Methods and practices. *IEEE TPAMI* **44**(5), 2548–2566 (2020)
10. Jing, W., Wang, J., Di, D., Li, D., Song, Y., Fan, L.: Multi-modal hypergraph contrastive learning for medical image segmentation. *Pattern Recognition* **165**, 111544 (2025)
11. Kalb, T., Mauthe, B., Beyerer, J.: Improving replay-based continual semantic segmentation with smart data selection. In: 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC). pp. 1114–1121. IEEE (2022)
12. Liu, H., Wei, D., Lu, D., Sun, J., Wang, L., Zheng, Y.: M3ae: multimodal representation learning for brain tumor segmentation with missing modalities. In: AAAI. vol. 37, pp. 1657–1665 (2023)
13. Rees, J.: Advances in magnetic resonance imaging of brain tumours. *Current opinion in neurology* **16**(6), 643–650 (2003)
14. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241. Springer (2015)
15. Ross, T.Y., Dollár, G.: Focal loss for dense object detection. In: CVPR. pp. 2980–2988 (2017)
16. Sawlani, V., Patel, M.D., Davies, N., Flintham, R., Wesolowski, R., Ughratdar, I., Pohl, U., Nagaraju, S., Petrik, V., Kay, A., et al.: Multiparametric mri: practical approach and pictorial review of a useful tool in the evaluation of brain tumours and tumour-like lesions. *Insights into imaging* **11**, 1–19 (2020)
17. Shi, J., Shang, C., Sun, Z., Yu, L., Yang, X., Yan, Z.: Passion: Towards effective incomplete multi-modal medical image segmentation with imbalanced missing rates. In: ACM International Conference on Multimedia. pp. 456–465 (2024)
18. Shi, J., Yu, L., Cheng, Q., Yang, X., Cheng, K.T., Yan, Z.: M2ftrans: Modality-masked fusion transformer for incomplete multi-modality brain tumor segmentation. *IEEE Journal of Biomedical and Health Informatics* (2023)
19. Tang, Q., Fan, L., Pagnucco, M., Song, Y.: Prototype-based image prompting for weakly supervised histopathological image segmentation. In: CVPR. pp. 30271–30280 (2025)
20. Tversky, A.: Features of similarity. *Psychological review* **84**(4), 327 (1977)
21. Van de Ven, G.M., Tuytelaars, T., Tolias, A.S.: Three types of incremental learning. *Nature Machine Intelligence* **4**(12), 1185–1197 (2022)
22. Wang, Z., Hong, Y.: A2fseg: Adaptive multi-modal fusion network for medical image segmentation. In: MICCAI. pp. 673–681. Springer (2023)
23. Yi, K., Chen, J., Wang, Y.G., Zhou, B., Lio, P., Fan, Y., Hamann, J.: Approximate equivariance so (3) needlet convolution. In: Topological, Algebraic and Geometric Learning Workshops 2022. pp. 189–198. PMLR (2022)
24. Yuan, B., Zhao, D.: A survey on continual semantic segmentation: Theory, challenge, method and application. *IEEE TPAMI* (2024)
25. Zeng, Z., Peng, Z., Yang, X., Shen, W.: Missing as masking: Arbitrary cross-modal feature reconstruction for incomplete multimodal brain tumor segmentation. In: MICCAI. pp. 424–433. Springer (2024)
26. 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: MICCAI. pp. 107–117. Springer (2022)

27. Zhang, Z., Liu, X., Chen, Z., Zhang, Y., Yue, H., Ou, Y., Sun, X.: Anatomical consistency distillation and inconsistency synthesis for brain tumor segmentation with missing modalities. arXiv preprint arXiv:2408.13733 (2024)
28. Zhang, Z., Yang, G., Zhang, Y., Yue, H., Liu, A., Ou, Y., Gong, J., Sun, X.: Tmformer: Token merging transformer for brain tumor segmentation with missing modalities. In: AAAI. vol. 38, pp. 7414–7422 (2024)
29. Zhao, Z., Yang, H., Sun, J.: Modality-adaptive feature interaction for brain tumor segmentation with missing modalities. In: MICCAI. pp. 183–192. Springer (2022)
30. Zhou, T., Canu, S., Vera, P., Ruan, S.: Latent correlation representation learning for brain tumor segmentation with missing mri modalities. IEEE Transactions on Image Processing **30**, 4263–4274 (2021)