

# Virtual Nodes Guided Dynamic Graph Neural Network for Brain Tumor Segmentation with Missing Modalities

Sha Tao    Jiao Pan    Yu Guo    Chao Yao  
University of Science and Technology Beijing, Beijing, China

## Abstract

Multimodal magnetic resonance imaging (MRI) is crucial for brain tumor segmentation, with many methods leveraging its four key modalities to capture complementary information for effective sub-region analysis. However, the absence of several modalities is very common in practice, leading to severe performance degradation in existing full-modality segmentation methods. Limited by the structured data model, recent works often adopt a multi-stage training strategy for full-modality and missing-modality scenarios, which increases training costs and inadequately addresses the interference of miss. In this work, we propose a graph-based one-stage framework for robust brain tumor segmentation with missing modalities. Specifically, we introduce modality-specific virtual nodes that serve as supplementary information sources to compensate for missing modalities. To enhance model robustness against arbitrary modality combinations, we leverage the inherent flexibility of graph networks to devise a dynamic connection strategy. This mechanism dynamically adjusts the adjacency matrix based on modality availability, preserving beneficial information flow while mitigating interference effects caused by missing modalities. Furthermore, we introduce heterogeneous weight matrices to improve the graph network’s adaptability to multimodal scenarios. Extensive experiments on the BRATS-2018 and BRATS-2020 datasets demonstrate that our method outperforms the state-of-the-art methods on almost all subsets of incomplete modalities.

## 1. Introduction

Multi-parametric magnetic resonance imaging (MRI) is the current standard for the quantitative evaluation of brain tumors in clinical practice [13]. Specifically, four MRI modalities are commonly used: T1-weighted (T1), contrast-enhanced T1-weighted (T1c), T2-weighted (T2), and Fluid Attenuated Inversion Recovery (FLAIR) images. Each modality delivers specific and valuable features that collectively improve the precision of brain tumor segmenta-

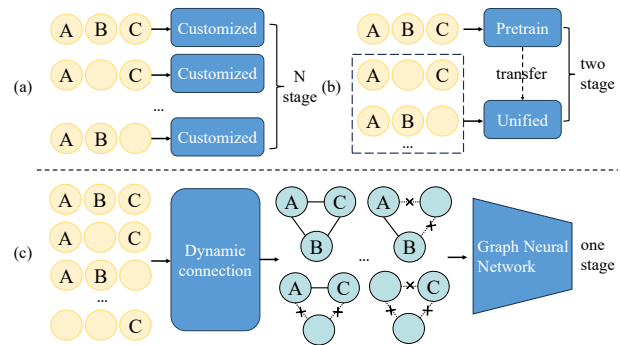


Figure 1. Comparison between our method and recent approaches. (a) Customized methods. (b) Recent unified methods. (c) Our graph-based approach.

tion. In recent years, plenty of works have explored the use of deep learning methods for brain tumor segmentation [14, 15]. These methods were optimized for the full-modality scenario and had already achieved excellent results. However, it is not always possible to acquire a complete set of MRIs due to practical constraints, such as image corruption, acquisition protocols, or unsuitable conditions of patients, and in more practical settings, some modalities may even be unavailable during training [29]. In such cases, these advanced methods often suffer from severe performance degradation.

To address the challenge of missing modalities, numerous methods have been proposed. A straightforward solution involves training a customized model for each possible subset of modalities [30, 35], as shown in Fig. 1 (a). While this strategy can yield promising results, it suffers from prohibitively high training and deployment costs, which grow exponentially with the number of modalities. This fundamental limitation severely restricts their applicability in real-world clinical settings. As a result, recent research has shifted toward learning a unified model capable of handling incomplete multimodal inputs, primarily by employing random modality dropout during training [10, 36]. However, structured data models (e.g., CNNs or Transformers) heavily rely on complete contextual information and struggle to

accommodate all modality combinations within a single-stage training framework. To overcome this limitation, many existing methods adopt a two-stage training strategy: either employing knowledge distillation [2, 4, 12, 24], where a full-modality model guides the learning of incomplete inputs, or training a generative model to reconstruct the missing modalities [19, 28].

However, these two-stage solutions ultimately stem from the rigidity of structured architectures such as CNNs and Transformers, whose computation paths are predefined for a fixed set of modalities and thus rely on complete cross-modality correspondence. As shown in Fig. 1(a–b), this fixed connectivity makes both customized and unified structured models sensitive to missing-modality patterns, since the expected feature interactions become ill-defined once one or more modalities are absent. In contrast, a graph formulation operates on the actually observed modality set: each modality is treated as a node, and missing modalities simply correspond to absent nodes or edges. As illustrated in Fig. 1(c), such adaptive connectivity enables message passing only among available modalities, making the model more robust under incomplete multimodal settings.

In this work, we propose a plug-and-play graph-based one-stage learning framework. Specifically, features extracted from modality-specific encoders are treated as nodes, which are then processed by a Graph Attention Network (GAT) for feature representation and fusion [27]. To enhance the GAT’s suitability for multimodal learning, we introduce a heterogeneous weight matrix that allows for modality-aware attention. To mitigate information loss due to missing modalities, we incorporate zero-initialized virtual nodes to capture modality-invariant features, thereby providing supplementary information even in the total absence of certain modalities. Furthermore, for the overall multimodal graph composed of both basic and virtual nodes, we design an adaptive edge connectivity strategy that supports diverse modality combinations, where information flows unidirectionally from existing nodes to those representing missing modalities. Experimenting on the public BraTS-2018 and BraTS-2020 datasets demonstrates that our method outperforms the state-of-the-art methods on almost all subsets of incomplete modalities. Even in the most extreme single-modality scenarios, our method still maintains robust performance.

Our key contributions can be summarized as follows:

- We propose a plug-and-play graph-based framework to handle missing modalities in brain tumor segmentation, which can be easily integrated into existing systems. The unstructured nature of graph architectures enables one-stage training directly on incomplete data, eliminating the need for additional supervision under full-modality settings.
- We introduce zero-initialized virtual nodes to mitigate

information loss caused by missing modalities. By enabling collaborative representation between base and virtual nodes, our model preserves comprehensive contextual information across arbitrary modality combinations.

- We design a dynamic edge connection mechanism where information flows unidirectionally from existing nodes to those representing missing modalities. This allows the model to flexibly adapt to varying combinations of available modalities within a single unified architecture.

## 2. Related Work

### 2.1. Multimodal Brain Tumor Segmentation with Missing Modalities

Existing methods for brain tumor segmentation with missing modalities can be broadly categorized into explicit and implicit approaches according to how missing modalities are represented.

Explicit (reconstruction-based) methods recover missing modality information either at the image level or the feature level. Image-level methods synthesize missing modalities for full-modality segmentation (e.g., using GANs) [8, 18, 32–34], yet accurately reconstructing critical missing modalities from the remaining ones is often difficult or even infeasible. Feature-level methods instead infer representative embeddings for missing modalities. For example, Shen and Gao [26] used adversarial learning to align feature distributions, while Wang et al. [28] proposed a Shared-Specific Feature Modeling framework for modality-invariant and modality-specific representations. Although such methods promote cross-modality alignment, they may suppress modality-specific cues, which is detrimental when essential modality information is absent. In addition, both image- and feature-level approaches usually require auxiliary generative networks, increasing model complexity and computational cost.

Implicit methods focus on learning a multimodal latent feature space. Hetero-modal variational encoder-decoder [6] incorporated multimodal variational auto-encoders to reconstruct modalities from a shared latent variable, while LS3M [38] employs learnable sorting and state space modeling for end-to-end segmentation under arbitrary modality combinations. Many methods also leverage knowledge distillation [11], where teacher (full-modality) and student (missing-modality) models are trained to transfer latent knowledge [2, 4, 12, 24]. Compared with explicit reconstruction, implicit methods are easier to train, but often provide limited compensation for missing modalities and still rely on two-stage training.

Overall, most existing methods still rely on structured architectures and often require two-stage optimization to cope with missing modalities. In contrast, our method adopts a graph-based formulation that combines explicit missing-




Figure 2. Overview of the proposed framework. The detailed dynamic edge connection is shown in Fig. 3 and Algorithm 1.

modality representation with implicit feature learning.

## 2.2. Graph Neural Networks for Medical Image Segmentation

Graph Neural Networks (GNNs) [9] have previously been applied to medical image segmentation tasks. Yan et al. [31] first used the SLIC algorithm [1] to cluster MRIs into supervoxels, and then predicted the tissue type of each supervoxel. Inspired by this work, subsequent research has generally followed a similar workflow [7, 21, 25]. However, such methods typically depend on superpixel-based representations, which hinder the preservation of modality-specific characteristics and reduce the effectiveness of cross-modal information fusion. These limitations are especially critical in scenarios with missing modalities.

In addition, several prior works have incorporated graph-inspired mechanisms to tackle missing modalities. Yang et al. [32] introduced a graph-attention fusion module for multimodal MR image synthesis and tumor segmentation, and Zhao et al. [39] proposed a modality-adaptive feature interaction framework based on graph representation concepts. However, these methods mainly incorporate graph-inspired operations into otherwise structured architectures, and do not explicitly address how missing modalities should be represented or how information should propagate under different missing-modality patterns. By contrast, our work formulates missing-modality interaction through a plug-and-play graph interface, where virtual nodes are introduced to account for absent modalities and the graph connectivity is dynamically adapted to the observed modality subset. As such, the proposed framework differs from prior graph-inspired methods not only in architectural form, but also in how modality absence is modeled and how cross-modality information flow is defined.

## 3. Method

An overview of the proposed framework is shown in Fig. 2. We first extract modality-specific features from the multi-modal MRI inputs using dedicated encoders. In parallel, we construct a graph containing both image-derived nodes and zero-initialized virtual nodes, together with a default adjacency matrix defined for the complete-modality setting. The extracted features are then mapped to their corresponding nodes and processed by a Graph Attention Network (GAT) equipped with a heterogeneous weight matrix to better model multimodal relationships. To support arbitrary combinations of available modalities, we adopt random modality dropout, replacing missing-modality nodes with zero vectors and dynamically updating graph connectivity via our adaptive edge strategy. Finally, the enriched node representations are fed into the decoder to produce the segmentation map.

The main difference between training and inference lies in the availability of modalities. During training, all modalities are accessible, and missing-modality scenarios are simulated via random modality dropout. Below, we will provide a detailed explanation of each module.

### 3.1. Node Construction

We denote the complete set of modalities by  $M = \{\text{FLAIR}, \text{T1c}, \text{T1}, \text{T2}\}$ . Given an input image of  $\mathbf{x}^m \in \mathbb{R}^{1 \times D \times H \times W}$ , where  $W, H, D$  are the width, height, depth of the image, and  $m \in M$ , we utilize the specific convolutional encoder to generate feature maps. In this condition, the original images of all modalities are available, so we can get independent and complete features of each modality. Assume that the features obtained through the encoder are denoted as  $\mathbf{f}^m \in \mathbb{R}^{C \times D' \times H' \times W'}$ , where  $C, D', H'$  and  $W'$  refer to the number of channels and the corresponding feature dimensions obtained after several convolutional down-sampling operations. Through the transformation of feature dimensions, we convert them into nodes  $\mathbf{v}^m \in \mathbb{R}^{C \times F}$ ,




Figure 3. Dynamic edge connection, illustrating the simulation from the full-modality scenario to the missing. The dashed shapes represent virtual nodes; the hollow shapes represent missing nodes; and ‘x’ denotes a broken connection.

where  $F = D' \times H' \times W'$ .

To mitigate the information loss introduced by missing modalities, we introduce a set of learnable, zero-initialized virtual nodes  $\mathbf{p}^m \in \mathbb{R}^{C_p \times F}$  for each modality. These nodes act as modality-agnostic anchors that interact with the available modality nodes through message passing, allowing the model to preserve a stable and coherent representation even when some modalities are absent.

Finally, the missing-aware virtual nodes are attached to the basic nodes  $\mathbf{v}^m$  along with the input-length dimension to form extended nodes  $\mathbf{v}_p^m$ :

$$\mathbf{v}_p^m = [\mathbf{v}^m; \mathbf{p}^m] \quad (1)$$

where  $[\cdot; \cdot]$  represents the concatenation operation,  $\mathbf{v}_p^m \in \mathbb{R}^{P \times F}$ , and  $P = C + C_p$ . For a complete multimodal graph network, there are a total of  $4 \times P$  nodes. We refer to the  $i$ -th node of the  $m$ -th modality as  $\mathbf{v}_i^m$ , and if the modality is not specifically indicated,  $m$  can be omitted.

### 3.2. Edge Construction and Dynamic Adaptability

Predefined or fixed connectivity becomes unreliable when modalities are missing, as absent inputs lead to noisy or invalid edges. Therefore, we construct dynamic connections that adjust the graph topology according to the available modalities. In this section, we detail our multimodal edge design. Since the Graph Attention Network (GAT) achieves weighted connections between different nodes through an attention mechanism, the construction of edge structures primarily serves to describe the connectivity between nodes with value 0 and 1. We use  $e_{ij}$  to represent an edge from  $\mathbf{v}_j$  to  $\mathbf{v}_i$ .

As shown in upper part of Fig. 3, the multimodal graph structure in a full-modal scenario constructs  $e_{ij}$  according to the following rules: intra-modal nodes are bidirectionally

connected, and so are corresponding nodes across different modalities. Virtual nodes are updated only by nodes of their own modality but participate in updating all other nodes.

During training, we randomly apply dropout to mask certain modality nodes, which dynamically alters the edge connections. Without loss of generality, suppose modality  $m$  is dropped (in practice, the number of dropped modalities is randomly selected between 0 and  $N - 1$ , where  $N$  is the total number of modalities). As shown in lower part of Fig. 3, for any basic node  $i$  within modality  $m$ , its connections to all other nodes  $j$  are cut, meaning  $e_{ji} = 0$  while  $e_{ij}$  remains unchanged. This approach allows us to reconstruct the missing modality features by utilizing other modality features and the invariant features of the the missing modality itself, while avoiding noise caused by reconstruction from interfering with other existing modalities. The detailed procedure is presented in Algorithm 1.

Fundamentally, our approach shares the same objective as customized methods—defining a specific model for each possible modality combination. However, the key distinction lies in our novel use of graph networks’ structural flexibility. Through dynamic connectivity, we implicitly encode  $N$  distinct model configurations within a single unified graph architecture. This enables our model to generalize across arbitrary modality combinations at inference time, achieving both adaptability and efficiency within a single-stage framework.

### 3.3. Heterogeneous Graph Attention Network

To calculate the weights when updating nodes, Graph Attention Network (GAT) performs self-attention on the nodes [27]. Since GAT was originally designed for single-modality data, it did not account for the differences between modalities in a multimodal scenario, resulting in the weight matrix  $\mathbf{W}$  being shared across all nodes. As shown in Fig. 2, in this work, we heterogenize the weight matrix  $\mathbf{W}$  to better accommodate the needs of multimodal features:

$$\beta_{ij} = a(\mathbf{W}^m \mathbf{v}_i^m, \mathbf{W}^n \mathbf{v}_j^n) \quad (2)$$

where  $m, n \in \{FLAIR, T1c, T1, T2\}$ ,  $a$  is an attentional mechanism and  $\beta_{ij}$  represents the attention coefficients between node  $\mathbf{v}_i$  and  $\mathbf{v}_j$ . To simulate the modality-missing scenario, during training, we apply a random dropout to mask the basic nodes of 0 to  $N - 1$  modalities by zero-filling them. Meanwhile, the edge structure is updated according to the rules mentioned earlier. When updating node  $\mathbf{v}_i$ , we calculate the relevance coefficient  $\beta_{ij}$  only for nodes  $\mathbf{v}_j$  where  $e_{ij} = 1$ , and get the weight coefficient of  $\mathbf{v}_j$  relative to  $\mathbf{v}_i$ :

$$\alpha_{ij} = \text{softmax}_j(\beta_{ij}) = \frac{\exp(\beta_{ij})}{\sum_{k \in \mathcal{N}_i} \exp(\beta_{ik})} \quad (3)$$

---

**Algorithm 1** Dynamic Edge Connection Method
 

---

```

1: Input:
2: Two nodes  $\mathbf{v}_i^m$  and  $\mathbf{v}_j^n$ , where  $m$  and  $n$  represent
   their respective modality; basic nodes set  $\mathcal{N}_b$ ; virtual
   nodes set  $\mathcal{N}_v$ ; total nodes set  $\mathcal{N}_t$ ; random missing node
    $\mathbf{v}_k$ .
3: Initial:
4: if  $m == n$  then
5:    $e_{ij} = e_{ji} = 1$ 
6: else
7:   if  $\mathbf{v}_i^m \in \mathcal{N}_b$  and  $\mathbf{v}_j^n \in \mathcal{N}_b$  then
8:     if  $i = j$  then
9:        $e_{ij} = e_{ji} = 1$ 
10:    else
11:       $e_{ij} = e_{ji} = 0$ 
12:    end if
13:  else if  $\mathbf{v}_i^m \in \mathcal{N}_b$  and  $\mathbf{v}_j^n \in \mathcal{N}_v$  then
14:     $e_{ij} = 1; e_{ji} = 0$ 
15:  else if  $\mathbf{v}_i^m \in \mathcal{N}_v$  and  $\mathbf{v}_j^n \in \mathcal{N}_b$  then
16:     $e_{ij} = 0; e_{ji} = 1$ 
17:  else
18:     $e_{ij} = e_{ji} = 0$ 
19:  end if
20: end if
21: Dynamic change:
22: for  $\mathbf{v}_i \in \mathcal{N}_t$  do
23:    $e_{ik} = 0$ 
24: end

```

---

where  $\mathcal{N}_i$  is the neighborhood of node  $i$  in the graph, i.e.  $\mathcal{N}_i = \{j \mid e_{ij} = 1\}$ . In practice, to handle the disconnected edges caused by missing modalities (i.e.,  $e_{ij} = 0$ ) and ensure gradient propagation, we did not simply remove them from the computation graph. Instead, we employed a softer approach by assigning a very small value to the corresponding attention weights. This approach reduces the interference from the missing modalities while maintaining differentiability.

In GAT’s experiments, the attention mechanism  $a$  is a single-layer feedforward neural network, parametrized by a weight vector  $\mathbf{a}$ , which is also shared between nodes. We heterogenized the  $\mathbf{a}$  as well, and the final weight coefficient can be expressed as:

$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(\mathbf{a}^m[\mathbf{W}^m \mathbf{v}_i^m \parallel \mathbf{W}^n \mathbf{v}_j^n]))}{\sum_{k \in \mathcal{N}_i} \exp(\text{LeakyReLU}(\mathbf{a}^m[\mathbf{W}^m \mathbf{v}_i^m \parallel \mathbf{W}^{m_k} \mathbf{v}_k^{m_k}]))} \quad (4)$$

Subsequently, similar to GAT, a multi-head attention mechanism is introduced to update the nodes:

$$\mathbf{v}'_i = \sigma \left( \frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \mathbf{W}_k^m \mathbf{v}_j^m \right) \quad (5)$$

where  $K$  is the number of independent attention mechanisms.

### 3.4. Overall Optimization:

Due to the inconsistency between the data distribution perceived by the specific encoder and the subsequent model, and to ensure the quality of the generated features, we introduced an additional specific decoder for optimization. Additionally, for the complete features updated by the GAT, a multimodal decoder is employed to perform segmentation prediction. To optimize virtual nodes, which represent modality-invariant features, some methods minimize the differences between modalities [28]. However, this manual intervention tends to reduce the richness of the generated information, with only shared features generated across modalities. Therefore, we adopt a self-learning approach, allowing the model to autonomously discover and enhance the modality-invariant features without imposing rigid constraints. This method leverages the natural alignment of features across modalities, encouraging the model to learn richer and more diverse representations.

In general, our method generates two types of segmentation maps during training, including  $\mathbf{y}^s$  of specific decoder and  $\mathbf{y}^m$  of multimodal decoder. Therefore, the overall segmentation loss is defined between the ground-truth label  $\mathbf{y}$  and the two estimated segmentation maps respectively:

$$\mathcal{L}_{\text{total}} = \sum_{i \in M} \mathcal{L}(\mathbf{y}_i^s, \mathbf{y}) + \mathcal{L}(\mathbf{y}^m, \mathbf{y}) \quad (6)$$

where  $M = \{\text{FLAIR}, \text{T1c}, \text{T1}, \text{T2}\}$  and  $\mathcal{L}$  is the Dice loss [23]. Note that the specific encoder is only used during the training phase, while during inference, the segmentation results are obtained solely from the multimodal decoder.

## 4. Experiments

### 4.1. Datasets and Evaluation Metrics

We evaluate our method on two widely used brain tumor segmentation benchmarks: BraTS 2018 and 2020 [22]. Both datasets provide multi-modal MRI scans for each subject, including FLAIR, T1, T1ce, and T2 sequences, along with expert-annotated tumor labels. The annotations consist of four categories: necrotic and non-enhancing tumor core (NET), enhancing tumor (ET), edema (ED), and background (BG). These categories are further grouped into three sub-regions for evaluation: (1) whole tumor (WT), including all tumor regions; (2) tumor core (TC), which comprises NET and ET; and (3) enhancing tumor (ET). The BraTS 2018 and 2020 datasets contain 285 and 369 annotated cases respectively. Following previous works, we split them into training/validation/testing subsets with ratios of 199/29/57 for BraTS 2018 [6] and 219/50/100 for BraTS 2020 [5, 37]. Following the BraTS challenge, we use the

| Modality |    |     |    |  | Enhancing tumor |      |                   |             |             | Tumor core |      |                   |             |             | Whole tumor |      |                   |             |             |
|----------|----|-----|----|--|-----------------|------|-------------------|-------------|-------------|------------|------|-------------------|-------------|-------------|-------------|------|-------------------|-------------|-------------|
| Fl       | T1 | T1c | T2 |  | HVED            | mmF  | M <sup>3</sup> AE | MMC         | Ours        | HVED       | mmF  | M <sup>3</sup> AE | MMC         | Ours        | HVED        | mmf  | M <sup>3</sup> AE | MMC         | Ours        |
| •        | ○  | ○   | ○  |  | 23.8            | 39.3 | 35.6              | <b>49.6</b> | 46.0        | 57.9       | 61.2 | 66.4              | 70.0        | <b>74.9</b> | 84.4        | 86.1 | 88.7              | 86.2        | <b>89.0</b> |
| ○        | •  | ○   | ○  |  | 8.6             | 32.5 | 37.1              | 42.9        | <b>43.9</b> | 33.9       | 56.6 | 66.1              | 69.3        | <b>72.0</b> | 49.5        | 67.5 | 74.4              | 78.6        | <b>78.9</b> |
| ○        | ○  | •   | ○  |  | 57.6            | 72.6 | 73.7              | <b>79.0</b> | 78.0        | 59.6       | 75.4 | 82.9              | <b>86.6</b> | 84.1        | 53.6        | 72.2 | 75.8              | <b>80.4</b> | 76.6        |
| ○        | ○  | ○   | •  |  | 22.8            | 43.1 | 47.6              | <b>50.7</b> | 50.0        | 54.7       | 64.2 | 69.4              | 69.7        | <b>73.6</b> | 79.8        | 81.2 | 84.8              | 84.1        | <b>84.9</b> |
| •        | •  | ○   | ○  |  | 28.0            | 43.0 | 41.2              | 48.1        | <b>50.3</b> | 61.1       | 65.9 | 70.8              | 72.0        | <b>77.7</b> | 85.7        | 87.1 | 89.0              | 87.7        | <b>90.0</b> |
| •        | ○  | •   | ○  |  | 68.4            | 75.1 | 75.0              | 79.5        | <b>81.0</b> | 75.1       | 77.9 | 84.4              | 86.7        | <b>87.1</b> | 85.9        | 87.3 | 89.7              | 88.7        | <b>90.2</b> |
| •        | ○  | ○   | •  |  | 32.3            | 47.5 | 45.4              | 48.3        | <b>53.4</b> | 62.7       | 69.8 | 70.9              | 72.4        | <b>76.4</b> | 87.6        | 87.6 | 89.9              | 88.0        | <b>90.0</b> |
| ○        | •  | •   | ○  |  | 61.1            | 74.0 | 74.7              | 78.7        | <b>80.0</b> | 67.6       | 78.6 | 83.4              | <b>86.6</b> | 86.2        | 64.2        | 74.4 | 77.2              | 80.2        | <b>80.5</b> |
| ○        | •  | ○   | •  |  | 24.3            | 45.0 | 48.7              | 46.9        | <b>50.4</b> | 56.3       | 69.4 | 71.8              | 74.0        | <b>76.0</b> | 81.6        | 82.2 | 86.7              | 85.9        | <b>86.5</b> |
| ○        | ○  | •   | •  |  | 67.8            | 74.5 | 75.3              | 79.1        | <b>80.2</b> | 73.9       | 78.6 | 84.2              | 86.7        | <b>87.0</b> | 81.3        | 83.0 | 86.3              | <b>86.3</b> | 86.2        |
| •        | •  | •   | ○  |  | 68.6            | 75.5 | 74.0              | 78.1        | <b>81.3</b> | 77.1       | 79.8 | 84.1              | 86.6        | <b>87.3</b> | 86.7        | 87.3 | 88.9              | 88.3        | <b>90.3</b> |
| •        | •  | ○   | •  |  | 32.3            | 47.7 | 44.8              | 50.7        | <b>54.5</b> | 63.1       | 71.5 | 72.7              | 67.8        | <b>77.6</b> | 88.1        | 87.8 | 89.9              | 88.5        | <b>90.2</b> |
| •        | ○  | •   | •  |  | 68.9            | 75.7 | 73.8              | 79.1        | <b>80.7</b> | 76.8       | 79.6 | 84.6              | 86.6        | <b>87.4</b> | 88.1        | 88.1 | 90.2              | 88.9        | <b>90.3</b> |
| ○        | •  | •   | •  |  | 67.8            | 74.8 | 75.4              | 79.2        | <b>80.3</b> | 75.3       | 80.4 | 84.4              | 86.7        | <b>87.0</b> | 82.3        | 82.7 | 85.7              | <b>86.8</b> | 86.5        |
| •        | •  | •   | •  |  | 69.0            | 77.6 | 75.5              | 80.1        | <b>80.8</b> | 77.7       | 85.8 | 84.5              | <b>87.4</b> | 87.3        | 88.5        | 89.6 | 90.1              | 89.0        | <b>90.3</b> |
| Mean     |    |     |    |  | 46.8            | 59.9 | 59.9              | 64.7        | <b>66.0</b> | 64.8       | 73.0 | 77.4              | 79.3        | <b>81.4</b> | 79.2        | 82.9 | 85.8              | 85.8        | <b>86.7</b> |

Table 1. Brain tumor segmentation results for different modality combinations on BraTS 2018 dataset. Present and missing modalities are denoted by • and ○. The best results are bolded.

Dice similarity coefficient (DSC) for performance quantification.

## 4.2. Implementation Details

We adopt 3D U-Net [36] as the backbone, using its encoder as modality-specific feature extractors and reusing the decoder for both single-modality and multimodal decoding. The key difference lies in the number of channels: one for the specific decoder and four for the multimodal decoder. The framework is implemented in PyTorch 1.7 and trained on four NVIDIA 4090 GPUs with a batch size of 1 for 1000 epochs. We use the Adam optimizer [17] with an initial learning rate of 0.0002 and a cosine decay schedule [20]. During training, input volumes are randomly cropped from 240×240×155 to 128×128×128 voxels to reduce memory consumption. We preserve complete modality inputs for the modality-specific encoders, and simulate missing modalities by randomly zeroing out selected features when constructing graph nodes. At inference time, all  $2^N - 1$  modality combinations are evaluated by skipping the corresponding encoders and graph connections for missing modalities. For fair and consistent evaluation on full-resolution images, we adopt a sliding window strategy following [3].

## 4.3. Comparison with the State of the Art (SOTA)

The experimental results for brain tumor segmentation with missing modalities on BraTS2018 are shown in Tab 1, which compares our method with current state-of-the-art methods, including U-HVED [6], mmFormer [36], M<sup>3</sup>AE [19], MMCFormer [16]. Among them, MMCFormer is customized and the others are unified. Notably,

mmFormer serves as a representative one-stage method, whereas M<sup>3</sup>AE exemplifies the two-stage paradigm.

Compared with the latest state-of-the-art method MM-CFormer, our approach achieves superior performance on BraTS 2018, with Dice improvements of 1.3% on enhancing tumor (ET), 2.1% on tumor core (TC), and 0.9% on whole tumor (WT). While the average gains appear moderate, our method demonstrates consistent advantages across different modality combinations. Notably, when the most informative modality (T1ce) is missing (e.g., lines 5, 7, 9, and 12), we observe significant gains, with average improvements of 3.65% on ET and 4.3% on TC. This highlights the model’s capacity to compensate for missing critical information. Overall, our method achieves the best Dice scores in 12 out of 15 combinations for ET, TC, and WT, respectively. In the remaining cases, the performance gap is clinically negligible, with average differences below 1.5% compared to the top-performing baseline. Compared with other representative methods, our approach also achieves notable performance gains. It is worth noting that, compared with mmFormer, the primary difference lies in the mechanism for modeling multimodal interactions: mmFormer employs a Transformer between the encoder and decoder of a 3D U-Net, whereas our method adopts a graph neural network for the same purpose. Considering this key architectural distinction and the corresponding performance gains, the results suggest that graph-based modeling offers clear advantages for capturing modality relationships in missing-modality scenarios.

Moreover, as shown in Fig. 4, our method achieves a favorable trade-off between model efficiency and segmen-

| Modality |    |     |    | Enhancing tumor |      |                   |             |             | Tumor core |      |                   |             |             | Whole tumor |      |                   |             |             |
|----------|----|-----|----|-----------------|------|-------------------|-------------|-------------|------------|------|-------------------|-------------|-------------|-------------|------|-------------------|-------------|-------------|
| Fl       | T1 | T1c | T2 | HVED            | mmF  | M <sup>3</sup> AE | LS3M        | Ours        | HVED       | mmF  | M <sup>3</sup> AE | LS3M        | Ours        | HVED        | mmF  | M <sup>3</sup> AE | LS3M        | Ours        |
| ●        | ○  | ○   | ○  | 12.9            | 40.5 | 40.5              | 40.3        | <b>43.8</b> | 34.6       | 66.2 | 68.0              | <b>72.0</b> | 70.3        | 69.9        | 82.4 | 86.5              | 88.7        | <b>89.0</b> |
| ○        | ●  | ○   | ○  | 7.3             | 34.0 | 39.9              | 36.6        | <b>44.4</b> | 27.3       | 61.2 | 66.0              | 66.7        | <b>71.6</b> | 46.8        | 74.4 | 76.7              | 79.3        | <b>79.5</b> |
| ○        | ○  | ●   | ○  | 24.9            | 68.9 | 72.4              | 77.7        | <b>81.6</b> | 35.5       | 78.0 | 81.4              | 83.6        | <b>86.9</b> | 46.8        | 74.3 | 73.9              | 78.2        | <b>78.9</b> |
| ○        | ○  | ○   | ●  | 24.3            | 45.6 | 46.0              | 45.3        | <b>46.2</b> | 37.7       | 69.2 | 70.3              | <b>72.3</b> | 69.4        | 54.0        | 83.1 | 86.1              | <b>87.3</b> | 85.1        |
| ●        | ●  | ○   | ○  | 22.0            | 43.6 | 43.2              | 45.1        | <b>49.3</b> | 38.3       | 69.6 | 73.8              | 75.3        | <b>75.5</b> | 58.3        | 84.6 | 89.4              | 90.6        | <b>90.7</b> |
| ●        | ○  | ●   | ○  | 30.0            | 69.8 | 74.7              | 79.2        | <b>81.7</b> | 42.2       | 80.4 | 82.0              | 86.3        | <b>87.4</b> | 61.5        | 84.5 | 89.5              | 91.0        | <b>91.7</b> |
| ●        | ○  | ○   | ●  | 29.4            | 48.1 | 47.3              | <b>51.5</b> | 51.3        | 43.4       | 71.6 | 75.0              | <b>76.9</b> | 74.7        | 64.5        | 85.8 | 89.3              | 91.5        | <b>92.1</b> |
| ○        | ●  | ●   | ○  | 33.6            | 71.1 | 75.4              | 79.5        | <b>79.6</b> | 44.9       | 79.9 | 82.4              | <b>84.1</b> | 80.9        | 62.9        | 78.0 | 78.1              | 82.6        | <b>87.5</b> |
| ○        | ●  | ○   | ●  | 32.1            | 45.9 | 46.6              | 47.0        | <b>51.9</b> | 45.0       | 70.9 | 72.5              | 74.3        | <b>75.2</b> | 64.3        | 84.0 | 87.2              | <b>88.9</b> | 87.7        |
| ○        | ○  | ●   | ●  | 36.2            | 70.7 | 76.8              | 78.7        | <b>81.8</b> | 47.5       | 80.8 | 83.0              | 85.2        | <b>87.1</b> | 65.8        | 84.1 | 87.4              | <b>89.2</b> | 87.7        |
| ●        | ●  | ●   | ○  | 39.4            | 70.1 | 75.9              | <b>80.8</b> | 80.0        | 49.1       | 80.2 | 82.4              | 86.1        | <b>87.6</b> | 67.0        | 85.3 | 90.0              | 91.5        | <b>91.6</b> |
| ●        | ●  | ○   | ●  | 38.1            | 48.4 | 48.2              | 51.5        | <b>53.7</b> | 49.4       | 72.0 | 75.1              | <b>76.9</b> | 76.2        | 68.4        | 86.2 | 90.4              | 91.5        | <b>91.7</b> |
| ●        | ○  | ●   | ●  | 40.9            | 71.6 | 77.1              | 80.5        | <b>81.5</b> | 51.3       | 81.3 | 83.1              | 85.8        | <b>87.4</b> | 69.7        | 86.1 | 90.2              | 91.9        | <b>92.1</b> |
| ○        | ●  | ●   | ●  | 43.2            | 70.7 | 77.4              | <b>81.5</b> | 80.2        | 52.7       | 81.1 | 84.1              | 85.8        | <b>87.4</b> | 70.4        | 84.6 | 88.6              | <b>89.8</b> | 88.2        |
| ●        | ●  | ●   | ●  | 45.3            | 71.4 | 78.0              | <b>80.9</b> | 79.8        | 54.2       | 81.2 | 84.4              | 85.9        | <b>87.5</b> | 71.4        | 86.4 | 90.6              | 92.2        | <b>92.3</b> |
| Mean     |    |     |    | 30.6            | 58.0 | 61.3              | 63.6        | <b>65.8</b> | 43.5       | 74.9 | 77.6              | 79.8        | <b>80.8</b> | 62.8        | 82.9 | 86.3              | 88.2        | <b>88.4</b> |

Table 2. Brain tumor segmentation results for different modality combinations on BraTS2020 dataset. Present and missing modalities are denoted by ● and ○. The best results are bolded.

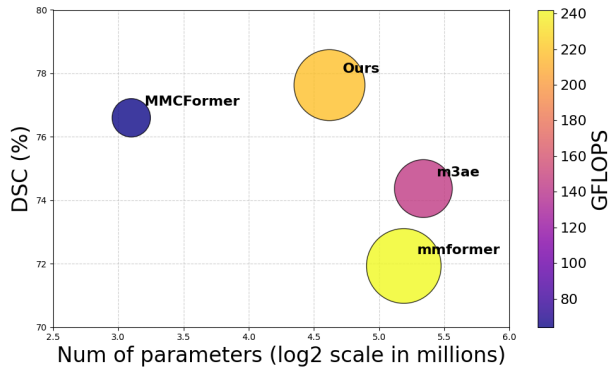


Figure 4. Deployment model size (in log<sub>2</sub> scale) and mean Dice similarity scores (DSCs) across all missing-modal situations on BraTS 2018 test split; circle size indicates GFLOPS.

tation performance. It operates with significantly fewer parameters than existing state-of-the-art methods, while maintaining superior accuracy. This efficiency is especially valuable in real-world scenarios with limited computational resources. Although MMCFormer is lightweight per model, it requires training a separate model for each missing modality setting, resulting in exponential parameter growth as the number of modalities increases. In contrast, our method supports arbitrary modality combinations within a single unified model, offering both robustness and scalability.

To evaluate the generalizability of our method, we conducted experiments on the BraTS 2020 dataset and compared it with state-of-the-art methods including U-HVED, mmFormer, M<sup>3</sup>AE, and LS3M [38]. As shown in Tab. 2,




Figure 5. Example segmentation results of our method on BraTS2018, focusing on four extreme cases where only one single modality remains, as well as the case with all modalities available. Green: peritumoral edema; yellow: enhancing tumor; and red: necrotic and non-enhancing tumor core.

our method consistently outperforms all baselines across the three tumor subregions, achieving average Dice improvements of 2.2% (ET), 1.0% (TC), and 0.2% (WT) over the strongest baseline LS3M. It achieves the highest Dice scores in 11/15 ET, 10/15 TC, and 11/15 WT cases.

Fig. 5 presents qualitative results under extreme conditions, where only a single modality is provided as input. These visualizations demonstrate that our model can still produce accurate and reliable segmentations, highlighting its robustness to missing modalities. In particular, results with only the T1c modality show clearly superior delineation of the enhancing tumor region, reaffirming the importance of T1c in capturing tumor enhancement.

Figure 6 presents qualitative comparisons between our method and two state-of-the-art approaches, mmFormer [36] and M3AE [19], under different modality combina-

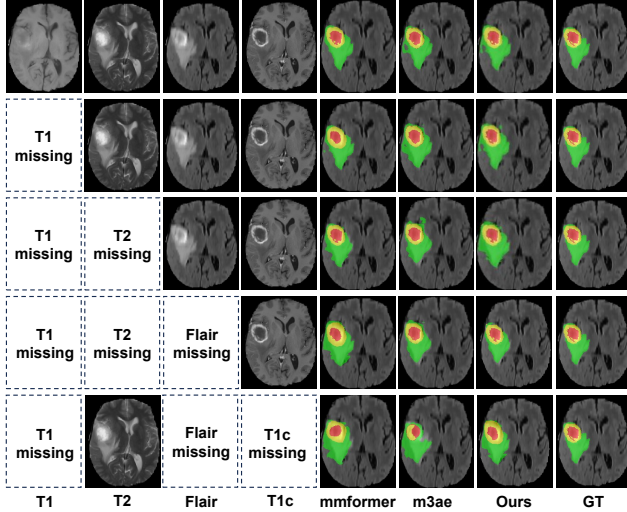


Figure 6. Qualitative segmentation results of different methods, including mmFormer, M3AE, and our approach, under various missing-modality settings. Green: peritumoral edema; yellow: enhancing tumor; and red: necrotic and non-enhancing tumor core.

tions. Consistent with the quantitative results, all methods achieve reasonable segmentation quality when the most informative modality, T1c, is available, even in the single-modality setting. However, when only T2 is available, both mmFormer and M3AE show clear performance degradation, particularly in the red and yellow regions, with evident spatial misalignment. In contrast, our method remains more robust and produces more coherent delineations across all tumor subregions. These results demonstrate that the proposed virtual nodes guided dynamic graph effectively preserves structural information under missing-modality conditions, highlighting its advantage in extreme missing-modality scenarios.

#### 4.4. Ablation Study

To evaluate the contribution of each component in our framework, we conducted ablation studies summarized in Table 3. Specifically, we removed one component at a time—virtual nodes, dynamic edge connections, or the heterogeneous weight matrix—to assess its individual impact. For the variant without dynamic connections, we adopted a fully connected static graph as a baseline. For the variant without the heterogeneous matrix, we used a standard Graph Attention Network (GAT) formulation.

The results show that removing any single component degrades performance, with the most notable drop observed in enhancing tumor segmentation. This is consistent with clinical observations, as ET segmentation heavily relies on the presence of T1c. The tumor core and whole tumor regions are comparatively less sensitive, showing only moderate declines. Notably, virtual nodes play a key role in com-

|     | ablation method      | ET          | TC          | WT          |
|-----|----------------------|-------------|-------------|-------------|
| w/o | virtual nodes        | 62.3        | 79.4        | 86.2        |
| w/o | dynamic connection   | 64.6        | 80.5        | 86.3        |
| w/o | heterogeneous matrix | 62.9        | 80.9        | 86.3        |
|     | total method         | <b>66.0</b> | <b>81.4</b> | <b>86.7</b> |

Table 3. Ablation study on BraTS2018. The mean performance (DSC, %) over all modality combinations is reported.

| length | ET          | TC          | WT          | Avg         |
|--------|-------------|-------------|-------------|-------------|
| 0      | 62.3        | 79.4        | 86.2        | 75.9        |
| 16     | 65.4        | 80.9        | 86.6        | 77.6        |
| 32     | <b>66.0</b> | <b>81.4</b> | <b>86.7</b> | <b>78.0</b> |
| 64     | 63.3        | 79.6        | 86.6        | 76.5        |
| 128    | 63.1        | 78.4        | 86.5        | 76.0        |

Table 4. Different length of virtual nodes of our model experimented on BraTS2018. The mean performance (DSC, %) over all modality combinations is reported.

pensating for missing modalities, and their removal causes the largest performance degradation, particularly for the ET region.

Table 4 examines the impact of virtual node length. While intuition might suggest that a longer virtual node enhances performance by enabling richer feature representation from missing modalities, our results show optimal performance at a length of 32—beyond which performance declines significantly. This counterintuitive finding indicates inherent limits to the modality-invariant information inferable from missing inputs: a moderate length (32) suffices to capture essential features, whereas larger dimensions introduce redundant or noisy information that impairs generalization. Additionally, increasing the virtual node length elevates computational overhead in memory and inference time. Thus, we select a length of 32, as it balances representational capacity and computational cost to achieve the best trade-off between effectiveness and efficiency.

## 5. Conclusion

In this work, we present a graph-based framework for multimodal brain tumor segmentation that effectively addresses missing-modality scenarios. By introducing virtual nodes and an adaptive connectivity mechanism, our method accommodates arbitrary modality combinations within a unified one-stage training paradigm and avoids the inefficiency of conventional two-stage pipelines. Extensive experiments on BraTS2018 and BraTS2020 show that our approach achieves competitive or superior performance compared with state-of-the-art methods. In future work, we plan to extend the framework to broader medical segmentation tasks and explore its potential as a plug-and-play solution for missing-modality challenges beyond the medical domain.

## Acknowledgements

This work was supported by the National Key Research and Development Program of China (Grant No. 2022ZD0118001), the National Natural Science Foundation of China (Grant Nos. 62332017, U22A2022), and was conducted in part at the MOE Key Laboratory of Advanced Materials and Devices for Post-Moore Chips, the Beijing Key Laboratory of Big Data Innovation and Application for Skeletal Health Medical Care, and the Beijing Advanced Innovation Center for Materials Genome Engineering.

## References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2274–2282, 2012. 3
- [2] Reza Azad, Nika Khosravi, and Dorit Merhof. Smu-net: Style matching u-net for brain tumor segmentation with missing modalities. In *International conference on medical imaging with deep learning*, pages 48–62. PMLR, 2022. 2
- [3] Cheng Chen, Qi Dou, Yueming Jin, Hao Chen, Jing Qin, and Pheng-Ann Heng. Robust multimodal brain tumor segmentation via feature disentanglement and gated fusion. In *International conference on medical image computing and computer-assisted intervention*, pages 447–456. Springer, 2019. 6
- [4] Cheng Chen, Qi Dou, Yueming Jin, Quande Liu, and Pheng Ann Heng. Learning with privileged multimodal knowledge for unimodal segmentation. *IEEE transactions on medical imaging*, 41(3):621–632, 2021. 2
- [5] Yuhang Ding, Xin Yu, and Yi Yang. Rfnnet: Region-aware fusion network for incomplete multi-modal brain tumor segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3975–3984, 2021. 5
- [6] Reuben Dorent, Samuel Joutard, Marc Modat, Sébastien Ourselin, and Tom Vercauteren. Hetero-modal variational encoder-decoder for joint modality completion and segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 74–82. Springer, 2019. 2, 5, 6
- [7] Antonio Garcia-Uceda Juarez, Raghavendra Selvan, Zaigham Saghir, and Marleen de Bruijne. A joint 3d unet-graph neural network-based method for airway segmentation from chest cts. In *International workshop on machine learning in medical imaging*, pages 583–591. Springer, 2019. 3
- [8] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [9] Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE international joint conference on neural networks, 2005.*, pages 729–734. IEEE, 2005. 3
- [10] Mohammad Havaei, Nicolas Guizard, Nicolas Chapados, and Yoshua Bengio. Hemis: Hetero-modal image segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 469–477. Springer, 2016. 1
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [12] Minhao Hu, Matthis Maillard, Ya Zhang, Tommaso Ciceri, Giammarco La Barbera, Isabelle Bloch, and Pietro Gori. Knowledge distillation from multi-modal to mono-modal segmentation networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 772–781. Springer, 2020. 2
- [13] Michael Iv, Byung C Yoon, Jeremy J Heit, Nancy Fischbein, and Max Wintermark. Current clinical state of advanced magnetic resonance imaging for brain tumor diagnosis and follow up. In *Seminars in roentgenology*, pages 45–61. Elsevier, 2018. 1
- [14] Haozhe Jia, Weidong Cai, Heng Huang, and Yong Xia. H 2 nf-net for brain tumor segmentation using multimodal mr imaging: 2nd place solution to brats challenge 2020 segmentation task. In *International MICCAI Brainlesion Workshop*, pages 58–68. Springer, 2020. 1
- [15] Zeyu Jiang, Changxing Ding, Minfeng Liu, and Dacheng Tao. Two-stage cascaded u-net: 1st place solution to brats challenge 2019 segmentation task. In *International MICCAI brainlesion workshop*, pages 231–241. Springer, 2019. 1
- [16] Sanaz Karimijafarbigloo, Reza Azad, Amirhossein Kazerooni, Saeed Ebadollahi, and Dorit Merhof. Mmcformer: Missing modality compensation transformer for brain tumor segmentation. In *Medical imaging with deep learning*, pages 1144–1162. PMLR, 2024. 6
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [18] Dongwook Lee, Won-Jin Moon, and Jong Chul Ye. Assessing the importance of magnetic resonance contrasts using collaborative generative adversarial networks. *Nature Machine Intelligence*, 2(1):34–42, 2020. 2
- [19] Hong Liu, Dong Wei, Donghuan Lu, Jinghan Sun, Liansheng Wang, and Yefeng Zheng. M3ae: multimodal representation learning for brain tumor segmentation with missing modalities. In *Proceedings of the AAAI conference on artificial intelligence*, pages 1657–1665, 2023. 2, 6, 7
- [20] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 6
- [21] Qihang Ma, Siyuan Zhou, Chengye Li, Feng Liu, Yan Liu, Mingzheng Hou, and Yi Zhang. Dgrunit: Dual graph reasoning unit for brain tumor segmentation. *Computers in biology and medicine*, 149:106079, 2022. 3
- [22] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014. 5

- [23] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016. 5
- [24] Yansheng Qiu, Delin Chen, Hongdou Yao, Yongchao Xu, and Zheng Wang. Scratch each other’s back: Incomplete multi-modal brain tumor segmentation via category aware group self-support learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21317–21326, 2023. 2
- [25] Camillo Saueressig, Adam Berkley, Elliot Kang, Reshma Munbodh, and Ritambhara Singh. Exploring graph-based neural networks for automatic brain tumor segmentation. In *International Symposium: From Data to Models and Back*, pages 18–37. Springer, 2020. 3
- [26] Yan Shen and Mingchen Gao. Brain tumor segmentation on mri with missing modalities. In *International conference on information processing in medical imaging*, pages 417–428. Springer, 2019. 2
- [27] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. 2, 4
- [28] Hu Wang, Yuanhong Chen, Congbo Ma, Jodie Avery, Louise Hull, and Gustavo Carneiro. Multi-modal learning with missing modality via shared-specific feature modelling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15878–15887, 2023. 2, 5
- [29] Junze Wang, Lei Fan, Weipeng Jing, Donglin Di, Yang Song, Sidong Liu, and Cong Cong. Hypergraph tversky-aware domain incremental learning for brain tumor segmentation with missing modalities. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 283–293. Springer, 2025. 1
- [30] Yixin Wang, Yang Zhang, Yang Liu, Zihao Lin, Jiang Tian, Cheng Zhong, Zhongchao Shi, Jianping Fan, and Zhiqiang He. Acn: adversarial co-training network for brain tumor segmentation with missing modalities. In *International conference on medical image computing and computer-assisted intervention*, pages 410–420. Springer, 2021. 1
- [31] Zhang Yan, Kong Youyong, Wu Jiasong, Gouenou Coatrieux, and Shu Huazhong. Brain tissue segmentation based on graph convolutional networks. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 1470–1474. IEEE, 2019. 3
- [32] Heran Yang, Jian Sun, and Zongben Xu. Learning unified hyper-network for multi-modal mr image synthesis and tumor segmentation with missing modalities. *IEEE Transactions on Medical Imaging*, 42(12):3678–3689, 2023. 2, 3
- [33] Biting Yu, Luping Zhou, Lei Wang, Yinghuan Shi, Jurgen Fripp, and Pierrick Bourgeat. Ea-gans: edge-aware generative adversarial networks for cross-modality mr image synthesis. *IEEE transactions on medical imaging*, 38(7):1750–1762, 2019.
- [34] Ziqi Yu, Xiaoyang Han, Shengjie Zhang, Jianfeng Feng, Tingying Peng, and Xiao-Yong Zhang. Mousegan++: unpervised disentanglement and contrastive representation for multiple mri modalities synthesis and structural segmentation of mouse brain. *IEEE Transactions on Medical Imaging*, 42(4):1197–1209, 2022. 2
- [35] Yao Zhang, Jiawei Yang, Jiang Tian, Zhongchao Shi, Cheng Zhong, Yang Zhang, and Zhiqiang He. Modality-aware mutual learning for multi-modal medical image segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 589–599. Springer, 2021. 1
- [36] Yao Zhang, Nanjun He, Jiawei Yang, Yuexiang Li, Dong Wei, Yawen Huang, Yang Zhang, Zhiqiang He, and Yefeng Zheng. mmformer: Multimodal medical transformer for incomplete multimodal learning of brain tumor segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 107–117. Springer, 2022. 1, 6, 7
- [37] Zheyu Zhang, Gang Yang, Yueyi Zhang, Huanjing Yue, Aiping Liu, Yunwei Ou, Jian Gong, and Xiaoyan Sun. Tm-former: Token merging transformer for brain tumor segmentation with missing modalities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7414–7422, 2024. 5
- [38] Zheyu Zhang, Yayuan Lu, Feipeng Ma, Yueyi Zhang, Huanjing Yue, and Xiaoyan Sun. Incomplete multi-modal brain tumor segmentation via learnable sorting state space model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25982–25992, 2025. 2, 7
- [39] Zechen Zhao, Heran Yang, and Jian Sun. Modality-adaptive feature interaction for brain tumor segmentation with missing modalities. In *International conference on medical image computing and computer-assisted intervention*, pages 183–192. Springer, 2022. 3