

Multi-modal Learning with Missing Modality via Shared-Specific Feature Modelling

Hu Wang*, Yuanhong Chen*, Congbo Ma*, Jodie Avery*, Louise Hull*, Gustavo Carneiro[°]

*The University of Adelaide, Adelaide, Australia

[°]Centre for Vision, Speech and Signal Processing, University of Surrey, UK

*{firstname.lastname}@adelaide.edu.au , [°]g.carneiro@surrey.ac.uk

Abstract

The missing modality issue is critical but non-trivial to be solved by multi-modal models. Current methods aiming to handle the missing modality problem in multi-modal tasks, either deal with missing modalities only during evaluation or train separate models to handle specific missing modality settings. In addition, these models are designed for specific tasks, so for example, classification models are not easily adapted to segmentation tasks and vice versa. In this paper, we propose the Shared-Specific Feature Modelling (ShaSpec) method that is considerably simpler and more effective than competing approaches that address the issues above. ShaSpec is designed to take advantage of all available input modalities during training and evaluation by learning shared and specific features to better represent the input data. This is achieved from a strategy that relies on auxiliary tasks based on distribution alignment and domain classification, in addition to a residual feature fusion procedure. Also, the design simplicity of ShaSpec enables its easy adaptation to multiple tasks, such as classification and segmentation. Experiments are conducted on both medical image segmentation and computer vision classification, with results indicating that ShaSpec outperforms competing methods by a large margin. For instance, on BraTS2018, ShaSpec improves the SOTA by more than 3% for enhancing tumour, 5% for tumour core and 3% for whole tumour.¹

1. Introduction

Recently, multi-modal learning has attracted much attention by research and industry communities in both computer vision and medical image analysis. Audio, images and short videos are becoming common types of media being used for multiple types model prediction in many different appli-

cations, such as sound source localisation [6], self-driving vehicles [32] and vision-and-language applications [28,33]. Similarly, in medical domain, combining different modalities to improve diagnosis accuracy has become increasingly important [9,29]. For instance, Magnetic Resonance Imaging (MRI) is a common tool for brain tumour detection, which does not depend only on one type of MRI image, but on multiple modalities (i.e. Flair, T1, T1 contrast-enhanced and T2). However, the multi-modal methods above usually require the completeness of all modalities for training and evaluation, limiting their applicability in real-world with missing-modality challenges when subsets of modalities may be missing during training and testing.

Such challenge has motivated both computer vision [20] and medical image analysis [5, 8, 13, 25] communities to study missing-modality multi-modal approaches. Wang et al. [31] proposed an adversarial co-training network for missing modality brain tumour segmentation. They specifically introduced a “dedicated” training strategy defined by a series of independent models that are specifically trained to each missing situation. Another interesting point about all previous methods is that they have been specifically developed either for (computer vision) classification [20] or (medical imaging) segmentation [5, 8, 13, 25, 31], making their extension to multiple tasks challenging.

In this paper, we propose a multi-model learning with missing modality approach, called **Shared-Specific Feature Modelling (ShaSpec)**, which can handle missing modalities in both training and testing, as well as dedicated training and non-dedicated training². Also, compared with previously models, ShaSpec is designed with a considerably simpler and more effective architecture that explores well-understood auxiliary tasks (e.g., the distribution alignment and domain classification of multi-modal features), which enables ShaSpec to be easily adapted to classification and segmentation tasks. The main contributions are:

- An extremely simple yet effective multi-modal learn-

¹This work received funding from the Australian Government through the Medical Research Futures Fund: Primary Health Care Research Data Infrastructure Grant 2020 and from Endometriosis Australia. G.C. was supported by Australian Research Council through grant FT190100525.

²Non-dedicated training refers to train one model to handle different missing modality combinations.

ing with missing modality method, called Shared-Specific Feature Modelling (ShaSpec), which is based on modelling and fusing shared and specific features to deal with missing modality in training and evaluation and with dedicated and non-dedicated training;

- To the best of our knowledge, the proposed ShaSpec is the first missing modality multi-modal approach that can be easily adapted to both classification and segmentation tasks given the simplicity of its design.

Our results on computer vision classification and medical imaging segmentation benchmarks show that ShaSpec achieves state-of-the-art performance. Notably, compared with recently proposed competing approaches on BraTS2018, our model shows segmentation accuracy improvements of more than 3% for enhancing tumour, 5% for tumour core and 3% for whole tumour.

2. Related Work

2.1. Multi-modal Learning Models

Multi-modal learning has attracted increasing attention from the research community. In medical image analysis, Dou et al. [9] introduced the chilopod-shaped architecture that optimised through modality-dependent feature normalisation and a knowledge distillation objective. From a different viewpoint of combining the uncertainty measurement with multi-modal learning, pixel-wise coherency [22] has been used for multi-modal learning through the optimisation of low-rank covariance metrics. Han et al. [12] designed a trusted multi-view classifier by adopting Dirichlet distribution to model the multi-modal uncertainties and fusing features via Dempster’s rule. Wang et al. [29] introduced an uncertainty-aware multi-modal learning model through cross-modal random network prediction.

In computer vision, Wang et al. [30] combined the idea of channel exchanging and multi-modal learning to fuse features. On video/audio classification and retrieval tasks, Patrick et al. [23, 24] proposed a self-supervised learning method to train multi-modal models on extra data, which significantly improved model performance. Chen et al. [6] designed a model to improve video-and-sound source localisation accuracy by defining a trusted tri-map middle-ground. Jia et al. [15] proposed a model for multi-view learning by constraining the view-specific features to be orthogonal towards view-shared ones. Despite providing some improvement over previous approaches, such an orthogonality constraint is fairly strong and not well-motivated by classification/segmentation objectives, so it may hamper the model’s ability to learn semantically rich representations. The methods above achieve promising results under the completeness assumption of full modalities. However, in real-world scenarios a subset of modalities may

not be available during training and evaluation.

Feature disentanglement methods [17, 18] aim to model the factors of data variations by learning representations that are modular (each latent dimension denotes a generative factor) and informative (representation has all generative factors) properties. Although our learned shared and specific features are also designed to be modular and informative, we do not aim to perform image (or other input data) reconstruction (i.e., the ShaSpec does not have any generative model) since we only target the learning of optimal representations for the classification and segmentation tasks. This lack of generative requirements significantly simplifies the training of our ShaSpec. Also, Jia et al.’s [15] method that learns shared and specific features is based on an arbitrary orthogonality criterion to de-correlate the features, which may not be the optimal way of learning modality-specific and modality-robust features. We argue that an optimisation function learns shared features through multi-modal feature distribution alignment and specific features via modality classification will likely lead to better performance. An important note about the multi-modal approaches above is that they have been specifically designed for specific tasks, and their adaptation from classification to segmentation, or vice versa, is not straightforward and has not actually been implemented.

2.2. Addressing Missing Modality in Multi-modal Learning

To overcome the missing modality issue in multi-modal learning, many methods have been developed. In computer vision, Ma et al. [20] proposed the SMIL model to deal with missing modality with a meta-learning algorithm that reconstructs the features of the missing modality data. Yin et al. [34] aimed to learn a unified subspace for incomplete and unlabelled multi-view data. In medical image analysis, Havaei et al. [13] developed a model called HeMIS to handle missing modalities by adopting statistical features (mean and variance) for decoding. Dorent et al. [8] extended the HeMIS model with a multi-modal variational auto-encoder (MVAE) that produces pixel-wise classifications based on mean and variance features. Similarly, auto-encoder structures have been adopted to reconstruct the missing modalities in a unsupervised learning scheme [4, 26]. Multiple approaches [14, 25, 31] proposed the learning of missing modality features from full modality models to improve the embeddings. Zhang et al. [35] introduced a vision transformer architecture for multi-modal brain tumour segmentation that aims to fuse features from all modalities into a set of comprehensive features.

The aforementioned models are mainly focused on reconstructing the missing modalities/features or introducing sophisticated architectures to solve missing modality problem. However, all of them neglect an essential point to

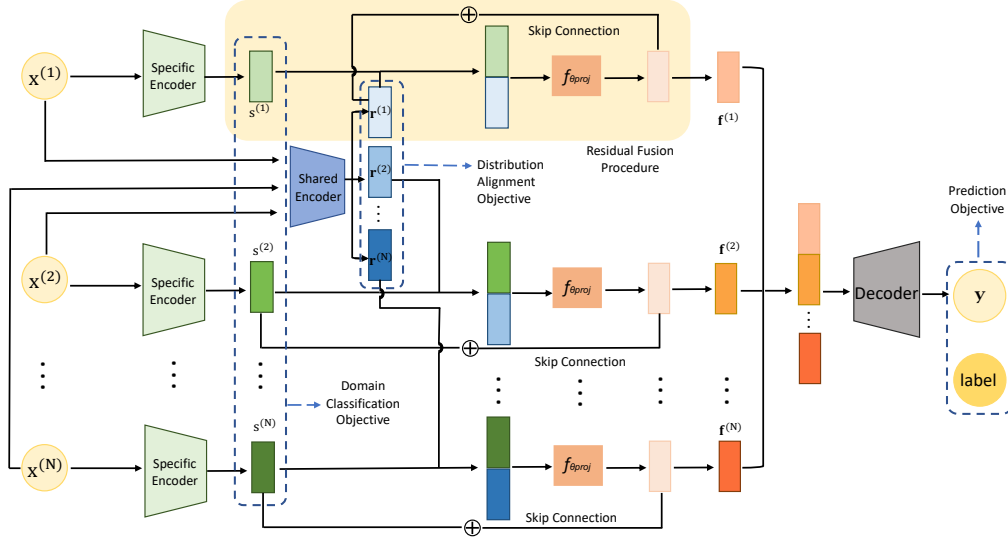


Figure 1. **Full-modality training and evaluation of ShaSpec.** All modalities $\{\mathbf{x}^{(i)}\}_{i=1}^N \in \mathcal{M}$ are passed through one shared encoder and individual specific encoders to produce the shared features $\{\mathbf{r}^{(i)}\}_{i=1}^N$ and specific features $\{\mathbf{s}^{(i)}\}_{i=1}^N$, respectively. Then, in a residual learning manner, the shared and specific features are fused with a linear projection $f_{\theta_{\text{proj}}}(\cdot)$ to get the fused features $\{\mathbf{f}^{(i)}\}_{i=1}^N$ for decoding. The dashed blue arrows indicate different objective functions.

tackle the missing modality challenge: how to learn the shared (i.e., modality-robust) and specific (i.e., modality-specific) features for optimisation of the model performance. Chen et al. [5] explored this direction and proposed a feature disentanglement and gated fusion model, called Robust-Mseg, for missing modality multi-modal segmentation. However, empirically, the algorithm faces unstable performance when tested on different missing modality scenarios, which we argue is caused by the high complexity of the model. Additionally, Robust-Mseg trained the appearance code just for the reconstruction of different modalities, which is weakly linked to the segmentation task. Moreover, from an implementation perspective³, the appearance code generation of the model proposed by [5] contains no missing modality drop and always requires full modality inputs, so it cannot handle missing data in training. Similarly to the multi-modal problem, when dealing with missing modality, methods are designed specifically for classification or segmentation and do not generalise to both tasks, which explains the lack of methods that can handle both tasks.

Our ShaSpec addresses the issues listed above with a model of rather simple but effective architecture. It achieves the goal by learning shared and specific features not only via the main task, but also with a distribution alignment and a domain classification tasks. The shared and specific features that contain rich information associated with the main task are finally fed into a decoder for prediction, which can be either classification or segmentation. Also,

³According to the official released code of [5] from <https://github.com/cchen-cc/Robust-Mseg>.

our model can handle missing modalities in training/testing and dedicated/non-dedicated missing modality training.

3. Methodology

3.1. Overall Architecture

Let us represent the N -modality data with $\mathcal{M}_j = \{\mathbf{x}_j^{(i)}\}_{i=1}^N$, where $\mathbf{x}_j^{(i)} \in \mathcal{X}$ denotes the j^{th} data sample and the superscript (i) indexes the modality. To simplify the notation, we omit the subscript j when that information is clear from the context. The label for each set \mathcal{M} is represented by $\mathbf{y} \in \mathcal{Y}$, where \mathcal{Y} represents the space for segmentation maps or classification categories. The model consists of a shared encoder denoted by $f_{\theta_{\text{sha}}} : \mathcal{X} \rightarrow \mathcal{R}$, specific encoders represented by $f_{\theta_{\text{spec}}}^{(i)} : \mathcal{X} \rightarrow \mathcal{S}$ for $i \in \{1, \dots, N\}$, a feature projection layer $f_{\theta_{\text{proj}}} : \mathcal{R} \times \mathcal{S} \rightarrow \mathcal{F}$ and a decoder $f_{\theta_{\text{dec}}} : \mathcal{F}^N \rightarrow \mathcal{Y}$. As shown in Fig. 1 and Fig. 2, the architecture suffers slight modifications while training and evaluating with full modalities or missing modalities. Please note that the ShaSpec model works for segmentation and classification, but the decoder in Fig. 1 and Fig. 2 only works for segmentation. For classification, the fused features are fed into fully connected (FC) layers. Below, we explain the evaluation and training processes with full and missing modalities. The methodology works for both dedicated and non-dedicated training.

3.2. Evaluation with Full and Missing Modalities

As shown in Fig. 1, in full modality training/evaluation, the process starts with the shared and specific branches run-

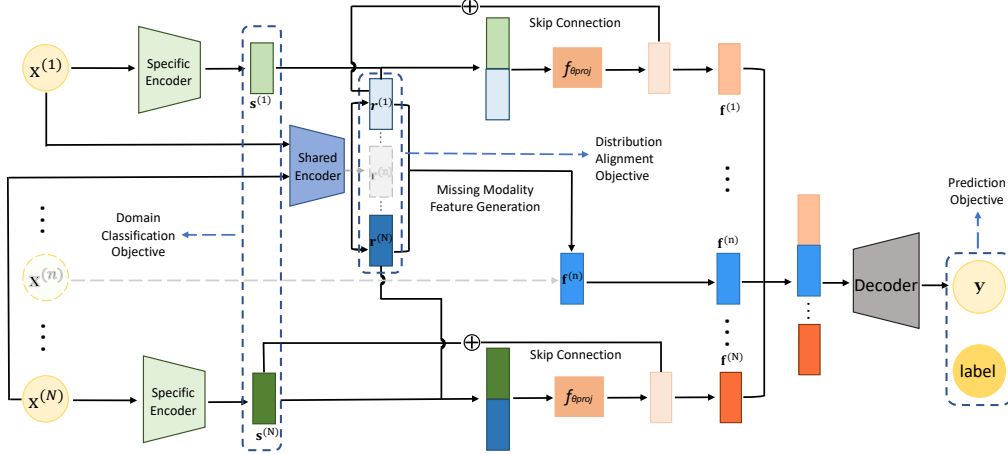


Figure 2. **Missing-modality training and evaluation of ShaSpec.** Without losing generality, we assume $\mathbf{x}^{(n)}$ is missing, where n can be $1, 2, \dots, N$. For available modalities $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n-1)}, \mathbf{x}^{(n+1)}, \dots, \mathbf{x}^{(N)}$, the shared-specific fused features $\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(n-1)}, \mathbf{f}^{(n+1)}, \dots, \mathbf{f}^{(N)}$ are extracted in the same way as in full modality. But for the missing modality data $\mathbf{x}^{(n)}$, the fused features $\mathbf{f}^{(n)}$ are generated from available shared features $\mathbf{r}^{(1)}, \dots, \mathbf{r}^{(n-1)}, \mathbf{r}^{(n+1)}, \dots, \mathbf{r}^{(N)}$ via a missing modality feature generation process. The dashed blue arrows indicate different objective functions.

ning in parallel, with

$$\mathbf{r}^{(i)} = f_{\theta^{\text{sha}}}(\mathbf{x}^{(i)}), \text{ and } \mathbf{s}^{(i)} = f_{\theta^{\text{spec}}}(\mathbf{x}^{(i)}), \quad (1)$$

for $i \in \{1, \dots, N\}$. Then, in residual fusion procedure, the shared and specific features are concatenated at the input of the projection layer, whose output is added as a residual to the shared features to form the semantically rich modality embedding, as follows:

$$\mathbf{f}^{(i)} = f_{\theta^{\text{proj}}}(\mathbf{r}^{(i)}, \mathbf{s}^{(i)}) + \mathbf{r}^{(i)}. \quad (2)$$

The final decoder then takes all embeddings to produce the output, with

$$\tilde{\mathbf{y}} = f_{\theta^{\text{dec}}}(\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(N)}), \quad (3)$$

where $\tilde{\mathbf{y}} \in \mathcal{Y}$. The intuition of the model above is that after extracting shared and specific features from each modality, the specific features $\{\mathbf{s}^{(i)}\}_{i=1}^N$ present modality-heterogeneous representations, while the shared features $\{\mathbf{r}^{(i)}\}_{i=1}^N$ capture the consistent features among modalities.

When a subset of the input modalities is missing, then the model architecture changes, as shown in Fig. 2. Without loss of generality, let us assume that the n^{th} modality $\mathbf{x}^{(n)} \in \mathcal{M}$ is missing. For all other available modalities that are not missing, i.e., $\{\mathbf{x}^{(i)}\}_{i=1, i \neq n}^N$, the process to extract $\{\mathbf{f}^{(i)}\}_{i=1, i \neq n}^N$ is the same as with the full modality in (1), but for the missing modality data, we directly generate the embedding $\mathbf{f}^{(n)}$ from the other available modalities via a missing modality feature generation process, defined by

$$\mathbf{f}^{(n)} = \frac{1}{N-1} \sum_{i=1, i \neq n}^N \mathbf{r}^{(i)}. \quad (4)$$

The model output is then produced with the decoder in Eq. (3). When we have more than one, but less than N missing modalities, we simply generate their features from the ones available using Eq. (4) and adapting the factor $\frac{1}{N-1}$ before summation.

3.3. Training with Full and Missing Modalities

For the model training, besides optimising for the main task (segmentation or classification), we introduce two auxiliary tasks, domain classification and distribution alignment, for the learning of the specific and shared feature representations, respectively.

3.3.1 Domain Classification Objective

Inspired by the domain adaptation technique from [10], we propose to adopt the domain classification objective (DCO) for the specific feature learning. The intuition is that if the specific features from a certain modality can be used to classify its domain (e.g., in brain tumour segmentation the domains can be Flair, T1, T1 contrast-enhanced or T2), then these specific features should contain valuable information that is specific for that modality. For domain classification, the cross-entropy (CE) loss is used for all available modalities. Formally, we have:

$$\ell_{dco}(\mathcal{D}, \theta^{\text{spec}}, \theta^{\text{dco}}) = - \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^N (\mathbf{t}^{(i)})^{\top} \log(f_{\theta^{\text{dco}}}(\mathbf{s}_j^{(i)})), \quad (5)$$

where $\mathbf{t}^{(i)} \in \{0, 1\}^N$ is the one-hot modality label with 1 at i^{th} position and 0 elsewhere, $\mathbf{s}_j^{(i)}$ denotes the i^{th} modality specific feature of the j^{th} training sample computed

from Eq. (1) (note that $i \in \{1, \dots, n-1, n+1, \dots, N\}$ if modality n is missing), $f_{\theta^{dco}} : \mathcal{S} \rightarrow \Delta^{N-1}$, with Δ^{N-1} denoting the probability simplex with N classes, and $\mathcal{D} = \{(\mathcal{M}_j, \mathcal{Y}_j)\}_{j=1}^{|\mathcal{D}|}$ is the training set.

3.3.2 Distribution Alignment Objective

The distribution alignment objective (DAO) is achieved by attempting to confuse the domain classifier by minimising the CE loss:

$$\ell_{dao}(\mathcal{D}, \theta^{\text{sha}}, \theta^{\text{dao}}) = - \sum_{j=1}^{|\mathcal{D}|} \sum_{i=1}^N (\mathbf{u}^{(i)})^\top \log(f_{\theta^{\text{dao}}}(\mathbf{r}_j^{(i)})), \quad (6)$$

where $\mathbf{u}^{(i)} = \frac{1}{N}$ is a uniform distribution for all modalities $i \in \{1, \dots, N\}$ (note that $i \in \{1, \dots, n-1, n+1, \dots, N\}$ if modality n is missing), and $f_{\theta^{\text{dao}}} : \mathcal{R} \rightarrow \Delta^{N-1}$ is the shared feature modality classifier. In Eq. (6), if the classification result can not distinguish the i^{th} modality from the others using the shared feature $\mathbf{r}_j^{(i)}$, then this is a robust shared feature representation.

Another option for this distribution alignment objective is the minimisation of the Kullback–Leibler divergence (KL divergence) between probabilities produced by the shared feature representations. To reduce the computation complexity, we project the feature onto a low dimensional space through a simple linear projection, as follows:

$$\ell_{dao}(\mathcal{D}, \theta^{\text{sha}}, \theta^{\text{dao}}) = \sum_{j=1}^{|\mathcal{D}|} \sum_{i,k=1}^N \text{KL}[\sigma(f_{\theta^{\text{dao}}}(\mathbf{r}_j^{(i)})), \sigma(f_{\theta^{\text{dao}}}(\mathbf{r}_j^{(k)}))] \quad (7)$$

where the $f_{\theta^{\text{dao}}}(\cdot)$ is the linear projection that produces an input for the softmax function $\sigma(\cdot)$, $KL(\cdot)$ is the Kullback–Leibler divergence operator. One more option for the DAO is the pairwise feature similarity, using

$$\ell_{dao}(\mathcal{D}, \theta^{\text{sha}}) = \sum_{j=1}^{|\mathcal{D}|} \sum_{i,k=1}^N \|\mathbf{r}_j^{(i)} - \mathbf{r}_j^{(k)}\|_p, \quad (8)$$

where $\|\cdot\|_p$ denotes the p-norm operator. In the ablation studies, we test different distribution alignment objectives.

3.3.3 Overall Objective

Besides the aforementioned DCO and DAO objectives, the objective for the main task is denoted by $\ell_{task}(\cdot)$ (e.g., cross-entropy loss for classification or Dice loss for segmentation). The overall objective to be minimised is:

$$\ell_{tot}(\mathcal{D}, \Theta) = \ell_{task}(\mathcal{D}, \theta^{\text{sha}}, \theta^{\text{spec}}, \theta^{\text{proj}}, \theta^{\text{dec}}) + \alpha \ell_{dao}(\mathcal{D}, \theta^{\text{sha}}, \theta^{\text{dao}}) + \beta \ell_{dco}(\mathcal{D}, \theta^{\text{spec}}, \theta^{\text{dco}}), \quad (9)$$

where $\Theta = \{\theta^{\text{sha}}, \theta^{\text{spec}}, \theta^{\text{proj}}, \theta^{\text{dao}}, \theta^{\text{dco}}, \theta^{\text{dec}}\}$; α and β are trade-off factors between different objective functions. In the ablation studies we test multiple values for α and β .

When dealing with missing modality n in training, the calculation of the missing modality feature $\mathbf{f}^{(n)}$ follows Eq. (4). This allows the optimisation of $\ell_{task}(\cdot)$. For the optimisation of $\ell_{dao}(\cdot)$ and $\ell_{dco}(\cdot)$, the losses from missing modality features $\mathbf{r}_j^{(n)}$ and $\mathbf{s}_j^{(n)}$ are omitted. Therefore, our proposed framework can seamlessly deal with missing modality issues in both training and evaluation.

4. Experiments

4.1. Datasets

We test our multi-modal learning with missing modality method on two datasets, the BraTS2018 for medical image segmentation and Audiovision-MNIST for computer vision classification. The **BraTS2018 Segmentation Challenge dataset** [1, 21] is used as a multi-modal learning with missing modality brain tumour sub-region segmentation benchmark, where the sub-regions are enhancing tumour (ET), tumour core (TC), and whole tumour (WT). BraTS2018 contains 3D multi-modal brain MRIs, including Flair, T1, T1 contrast-enhanced (T1c) and T2 with experienced imaging experts annotated ground-truth. It includes 285 cases for training (210 gliomas with high grade and 75 gliomas with low grade) and 66 cases for evaluation. The ground-truth of training set is publicly available, but the annotations of validation set is hidden and online evaluation⁴ is required.

Our missing modality experiments for computer vision classification are conducted on **Audiovision-MNIST dataset** [27]. Audiovision-MNIST is a multi-modal dataset consisting of 1500 samples of audio and image files. The images for digits 0 to 9 are of a size 28×28 that come from the MNIST dataset [16]. The audio dataset contains 1500 audio files that have been collected from the Free Spoken Digits Dataset⁵. For the representation of audio modality, mel-frequency cepstral coefficients (MFCCs) are adopted to transform each audio sample into size $20 \times 20 \times 1$. Following [20], we split the dataset into 70% for training and 30% for evaluation, according to the officially released code⁶.

4.2. Implementation Details

The ShaSpec model has a straightforward training process without much hyperparameter tuning. Implementation details on both datasets are described below.

BraTS2018: We adopted the 3D UNet (with 3D convolution and normalisation) as our backbone network, where the fusion of shared and specific features happens at the bottom of the UNet structure. A stochastic gradient descent optimizer with Nesterov momentum [3] of 0.99 is adopted for optimisation. The learning rate is set to 10^{-2} at the be-

⁴Online evaluation at <https://ipp.cbica.upenn.edu/>.

⁵The data information can be found at <https://github.com/Jakobovski/free-spoken-digit-dataset>.

⁶The code address is <https://github.com/mengmenm/SMIL>.

Modalities				Enhancing tumour					tumour Core					Whole tumour							
Fl	T1	T1c	T2	UHeMIS	UHVED	RbSeg	mmFm	ShaSpec	ShaSpec*	UHeMIS	UHVED	RbSeg	mmFm	ShaSpec	ShaSpec*	UHeMIS	UHVED	RbSeg	mmFm	ShaSpec	ShaSpec*
•	•	•	•	11.78	23.80	25.69	39.33	43.52	45.11	26.06	57.90	53.57	61.21	69.44	69.57	52.48	84.39	85.69	86.10	88.68	88.83
•	•	•	•	10.16	8.60	17.29	32.53	41.00	42.58	37.39	33.90	47.90	56.55	63.18	64.53	57.62	49.51	70.11	67.52	73.44	74.82
•	•	•	•	62.02	57.64	67.07	72.60	73.29	75.80	65.29	59.59	76.83	75.41	78.65	81.40	61.53	53.62	73.31	72.22	73.82	74.95
•	•	•	•	25.63	22.82	28.97	43.05	46.31	46.21	57.20	54.67	57.49	64.20	69.03	69.05	80.96	79.83	82.24	81.15	83.99	84.90
•	•	•	•	10.71	27.96	32.13	42.96	44.76	44.81	41.12	61.14	60.68	65.91	72.67	72.77	64.62	85.71	88.24	87.06	89.76	89.86
•	•	•	•	66.10	68.36	70.30	75.07	75.60	77.76	71.49	75.07	80.62	77.88	84.50	84.75	68.99	85.93	88.51	87.30	90.06	90.12
•	•	•	•	30.22	32.31	33.84	47.52	47.20	47.22	57.68	62.70	61.16	69.75	72.93	72.93	82.95	87.58	88.28	87.59	90.02	90.09
•	•	•	•	66.22	61.11	69.06	74.04	75.76	78.26	72.46	67.55	78.72	78.59	82.10	82.64	68.47	64.22	77.18	74.42	78.74	78.88
•	•	•	•	32.39	24.29	32.01	44.99	46.84	49.87	60.92	56.26	62.19	69.42	71.38	71.39	82.41	81.56	84.78	82.20	86.03	86.09
•	•	•	•	67.83	67.83	69.71	74.51	75.95	78.59	76.64	73.92	80.20	78.61	83.82	84.08	82.48	81.32	85.19	82.99	85.42	86.43
•	•	•	•	68.54	68.60	70.78	75.47	76.42	78.51	76.01	77.05	81.06	79.80	85.23	85.36	72.31	86.72	88.73	87.33	90.29	90.36
•	•	•	•	31.07	32.34	36.41	47.70	46.55	46.56	60.32	63.14	64.38	71.52	73.97	73.99	83.43	88.07	88.81	87.75	90.36	90.37
•	•	•	•	68.72	68.93	70.88	75.67	75.99	78.15	77.53	76.75	80.72	79.55	85.26	85.67	83.85	88.09	89.27	88.14	90.78	90.79
•	•	•	•	69.92	67.75	70.10	74.75	76.37	78.35	78.96	75.28	80.33	80.39	84.18	84.27	83.94	82.32	86.01	82.71	86.47	86.51
•	•	•	•	70.24	69.03	71.13	77.61	78.08	78.47	79.48	77.71	80.86	85.78	85.45	85.75	84.74	88.46	89.45	89.64	90.88	90.88
Average				46.10	46.76	51.02	59.85	61.58	63.08	62.57	64.84	69.78	72.97	77.45	77.88	74.05	79.16	84.39	82.94	85.92	86.26

Table 1. Model performance comparison of **segmentation** Dice score (normalised to 100%) on BraTS2018 of **non-dedicated training**. ShaSpec and ShaSpec* are the proposed models, with ShaSpec* being the model with prediction smoothness enhancement. The best and second best results for each column within a certain type of tumour are in **red** and **blue**, respectively.

Modalities				Enhancing tumour				tumour Core				Whole tumour			
Fl	T1	T1c	T2	KD-Net	ACN	ShaSpec	ShaSpec*	KD-Net	ACN	ShaSpec	ShaSpec*	KD-Net	ACN	ShaSpec	ShaSpec*
•	•	•	•	40.99	42.77	43.94	43.97	65.97	67.72	70.97	70.99	85.14	87.30	89.28	89.38
•	•	•	•	39.87	41.52	45.24	46.76	70.02	71.18	70.28	70.64	77.28	79.34	79.40	79.50
•	•	•	•	75.32	78.07	75.91	78.40	81.89	84.18	84.19	85.47	76.79	80.52	80.43	80.55
•	•	•	•	39.04	42.98	44.54	46.07	66.01	67.94	70.30	70.11	82.32	85.55	85.58	85.62
Average				48.81	51.34	52.41	53.80	70.97	72.76	73.92	74.30	80.38	83.18	83.67	83.76

Table 2. Model performance comparison of **segmentation** Dice score (normalised to 100%) on BraTS2018 of **dedicated training**.

ginning and decreased with cosine annealing strategy [19]. During the non-dedicated training of models, modalities are randomly dropped to simulate the modality-missing situations. For dedicated training of models, the missing modalities used for training are the same missing modalities in the evaluation. The ShaSpec model is trained for 180,000 iterations using all training data without model selection. We choose L1 loss as our distribution alignment objective and set $\alpha = 0.1$, $\beta = 0.02$ in Eq. (9). Then, we perform the official online evaluation using the segmentation masks produced by ShaSpec. We run ShaSpec with and without prediction smoothness enhancement to improve segmentation results, when we run our model with this enhancement. We labelled it as ‘‘ShaSpec*’’. This enhancement connects components within the surrounding voxels with two hops being considered as neighbours. Moreover, small regions with fewer voxels than a certain threshold are eliminated. By doing so, scattered small regions are cancelled, which in general leads to segmentation improvements.

Audiovision-MNIST: For the model training on Audiovision-MNIST dataset, we follow the SMIL paper [20] by dropping the sound modality data for a certain percentage and training all models for 60 epochs to keep a fair comparison. we adopted the image and sound encoders from SMIL, consisting of networks with a sequence of convolutional layers and fully connected (FC) layers with batch norm and dropout. For the rest of the ShaSpec architecture, after fusing the two modality features, 2 FC

layers with dropout are adopted for classification. The fusion of shared and specific features happens at the layer before the FC layers. Adam optimizer with 10^{-2} weight decay is used for model training. The initial learning rate is set to 10^{-3} which is decreased by 10% every 20 epochs.

The evaluation of model performance relies on the Dice score for BraTS2018 and classification accuracy for Audiovision-MNIST, where training and evaluation are performed on one 3090Ti NVIDIA Graphics Card.

4.3. Segmentation Results

The experimental results for the non-dedicated training (training the model once and evaluating it on different combinations of missing modalities) on BraTS2018 are shown in Tab. 1, which compares ShaSpec with current state-of-the-art (SOTA) methods, including U-HeMIS [13], U-HVED [8], Robust-MSeg (RbSeg) [5] and mmFormer (mmFm) [35]. Our proposed ShaSpec shows the best and second best results (the best in 45 out of 48 results) across almost all different combinations and tumour types, as shown in red and blue of Tab. 1. Moreover, the proposed ShaSpec outperforms the competing models by a large margin. For instance, when only T1 is available, ShaSpec surpasses the second best model (mmFormer) by 8.47% on enhancing tumour, 6.63% on tumour core and 5.92% on whole tumour. Similarly, when only T1, T1c and T2 are available, we observe improvements of 1.62% on enhancing tumour, 3.79% on tumour core and 3.76%

Audio rate	LowerB	UpperB	AutoEncoder	GAN	Full2miss	SMIL	ShaSpec
5%	92.35	98.22	89.78	89.11	90.00	92.89	93.33
10%	92.35	98.22	89.33	89.78	91.11	93.11	93.56
15%	92.35	98.22	89.78	88.67	92.23	93.33	93.78
20%	92.35	98.22	88.89	89.56	92.67	94.44	94.67

Table 3. Model performance comparison of classification accuracy of missing modality (by setting different available audio rates) on Audiovision-MNIST dataset. The lower bound (LowerB) is a LeNet [16] network trained with single modality (images only). The upper bound (UpperB) is a model trained with all data modalities (all images and audios). The best results for each row are bolded.

on whole tumour. The Prediction Smoothness Enhancement further boosts ShaSpec’s performance. On average, when compared with the second best competing methods, our model gets 3.23% performance gain for enhancing tumour, 4.91% for tumour core and 3.32% for whole tumour. Also, in Tab. 1 note that for the enhancing tumour segmentation, the prediction smoothness enhancement is influential, improving from 43.52% to 45.11% when only Flair is available, and from 75.76% to 78.26% when T1 and T1c are available. This may be caused by scattered segmentation masks predicted for enhancing tumour. Also, T1c contributes more than other modalities for enhancing tumour. For instance, when adding T1c as an available modality, the performance of the model increases greatly. For ShaSpec, we have 73.29% with T1c only vs. 43.52% with Flair only. Such observation resonates with the knowledge that enhancing tumour is clearly visible in T1c, but edema is not visible [5]. Similar results are shown in Tab. 2, where ShaSpec outperforms KD-Net [14] and ACN [31] for most of cases (we are best in 11 out of 12 cases). On average, our model surpasses the second best competing method by 2.46% on enhancing tumour, 1.54% on tumour core and 0.58% on whole tumour.

4.4. Classification Results

Following the SMIL setup [20], we train ShaSpec on both partial and full modality sub-datasets (images and audios). More specifically, full missing modality sub-dataset is formed by setting the audio modality rates in {5%, 10%, 15%, 20%}, which defines the proportion of the available audio data used for training. For this setup, the visual modality data is fully available. In the evaluation phase, only images are fed into the models. We compare our model with Auto-encoder [2], a model based on Generative Adversarial Networks [11], a method that distils multi-modality knowledge to train missing modality model [25], and SMIL [20]. A LeNet [16] network with single modality (images only) serves as the lower bound, and following SMIL, a model trained with full modalities (all images and audios) works as the upper bound. As shown in Tab. 3, ShaSpec performs well, especially under extreme missing modality (i.e., small audio rates), where our model achieves 93.33% accuracy compared with second best model with 92.89% on audio rate 5%. With the increasing of the au-

dio rate, the performance of all models improved, with ShaSpec still outperforming all other models. We argue that this is due to the outstanding ability of ShaSpec to extract information-rich shared-specific representations from all available modalities and the use of the shared representation to make up for the missing modalities.




Figure 3. Model performance with varying rates of missing modality data (for both image and audio) on Audiovision-MNIST.

To examine the missing modality representation learning of the proposed ShaSpec, we train models for varying rates of both image and audio (rates in 15%, 20%, 50%, 70% and 100%) and test with both modalities. The ShaSpec model is compared with SMIL in Fig. 3. For all settings, our model is always superior as shown by the bar chart.

4.5. Analyses

The selection of DAO loss function: As discussed in Sec. 3.3.2, DAO can rely on a CE loss, KL divergence loss or p-norm distance loss between prediction and ground-truth. The comparisons on BraTS2018 is presented in Tab. 4 for non-dedicated training, where only T1 is available for evaluation. It shows that the different choices of DAO perform similarly, but the L1 loss has the best results, followed by the KL divergence loss, while CE loss and MSE do not perform as well as the other losses.

Sensitivity of Eq. (9) to α and β : The results are shown in Fig. 4 by setting the α and β values to {0, 0.02, 0.1, 0.5, 0.7, 1}. When testing α values, we set β to 0.02; and when testing β , we set α to 0.1. When $\alpha = \beta = 1$, the results drop significantly, which can be explained by the scale of different loss functions, where too large weights

DAO Type	CE	KL	L1	MSE
Enhancing Tumour	41.23	40.41	42.58	43.19
Tumour Core	62.72	62.83	64.53	62.44
Whole Tumour	73.91	74.19	74.82	73.25

Table 4. Model ablation of different distribution alignment objectives for non-dedicated training, where only T1 is available for testing on BraTS2018.




Figure 4. Sensitivity of α and β for Eq. (9) for non-dedicated training, where only T1 is available for evaluation on BraTS2018.

for the auxiliary losses may disturb the gradient flow of the main task. In general, $\alpha = 0.1$ and $\beta = 0.02$ produces the best results. Small values for the weights of the auxiliary tasks contribute to the whole process, but do not interfere with the main task optimisation. Interestingly, when $\alpha = 0$ (only specific features are learned), the model can still segment the tumours to some extent by simple concatenation of specific features, which means that the specific features contain rich information. A similar conclusion can be reached when $\beta = 0$ (only shared features are learned).

Computational comparison: We compare ShaSpec with SMIL (using the official code released by Ma et al.) in terms of the number of model parameters, training/inference iteration time and GPU memory usage, where both models are trained/tested with batch-size of 4, and we estimate the average time consumption of 30 iterations on one 3090 GPU for a fair comparison. SMIL has 0.33M parameters, with training iterations and testing taking 0.1309s and 0.0019s, and during training and testing, the GPU memory usage started from 1430MiB, climbed to 24268MiB, and then casted an “out of memory” error in the end. On the other hand, the ShaSpec model has 0.22M parameters, takes 0.0257s for model training iteration and 0.0016s for model testing, and constantly consumes 1421MiB of GPU memory.

An additional classification experiment on X-ray + clinical texts: We conducted an extra classification experiment on OpenI [7]. We reorganised the OpenI dataset by only considering frontal images as visual inputs, and the ‘COMPARISON’ and ‘FINDINGS’ tags in the reports for textual inputs. Also, we reformulate the multi-label Chest-Xray classification into a binary classification problem (whether any chest problem exists or not). In total, the dataset contains 3851 pairs of visual-textual samples, which are split




Figure 5. t-SNE visualisation of shared and specific features of four modalities from all training data on BraTS2018. The shared features of four modalities are presented by ‘x’ in different colours, while the specific features of four modalities are presented by ‘o’ in different colours.

into 80% for training and 20% for evaluation. We adopt ResNet50 as visual backbone and an LSTM model with 128 hidden neurons as textual backbone. The single-modal ResNet50 obtained AUC=0.77 with image inputs only; and the single-modal LSTM has AUC=0.86 with text inputs only. The baseline multi-modal model with simple feature concatenation trained/evaluated with full modality reaches AUC=0.90. When trained with 30% missing image modality and evaluated only on texts, it achieves AUC=0.87. Our ShaSpec model shows better performance than the baseline model with AUC=0.89 (close to full modality).

Shared and Specific Feature Visualisation: To further display the effectiveness of ShaSpec, we visualise the shared and specific features with t-SNE in Fig. 5. Note that the shared features are clustered together, while the specific features are well-separated, which shows that both types of features are exactly what is expected from our algorithm.

5. Conclusion

In this paper, we propose the simple but effective ShaSpec method to address multi-modal learning with missing modalities in training/testing, for dedicated/non-dedicated training, and applied to segmentation and classification tasks. Empirically, it outperforms the state-of-the-art by a large margin in different tasks and settings because of the semantically rich shared and specific learned features that are strongly related to the main task. Through the t-SNE visualisation of the shared and specific feature space, we further verify the effectiveness of the method. In the future, we will test the proposed ShaSpec method on other tasks (e.g., regression) and datasets to further verify its generality and effectiveness.

References

- [1] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, 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. 5
- [2] Pierre Baldi. Autoencoders, unsupervised learning, and deep architectures. In *Proceedings of ICML workshop on unsupervised and transfer learning*, pages 37–49. JMLR Workshop and Conference Proceedings, 2012. 7
- [3] Aleksandar Botev, Guy Lever, and David Barber. Nesterov’s accelerated gradient and momentum as approximations to regularised update descent. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 1899–1903. IEEE, 2017. 5
- [4] Agisilaos Chatsias, Thomas Joyce, Mario Valerio Giuffrida, and Sotirios A Tsaftaris. Multimodal mr synthesis via modality-invariant latent representation. *IEEE transactions on medical imaging*, 37(3):803–814, 2017. 2
- [5] 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. 1, 3, 6, 7
- [6] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Localizing visual sounds the hard way. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16867–16876, 2021. 1, 2
- [7] Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310, 2016. 8
- [8] 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. 1, 2, 6
- [9] Qi Dou, Quande Liu, Pheng Ann Heng, and Ben Glocker. Unpaired multi-modal segmentation via knowledge distillation. In *IEEE Transactions on Medical Imaging*, 2020. 1, 2
- [10] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 4
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 7
- [12] Zongbo Han, Changqing Zhang, Huazhu Fu, and Joey Tianyi Zhou. Trusted multi-view classification. *arXiv preprint arXiv:2102.02051*, 2021. 2
- [13] 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, 2, 6
- [14] 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, 7
- [15] Xiaodong Jia, Xiao-Yuan Jing, Xiaoke Zhu, Songcan Chen, Bo Du, Ziyun Cai, Zhenyu He, and Dong Yue. Semi-supervised multi-view deep discriminant representation learning. *IEEE transactions on pattern analysis and machine intelligence*, 43(7):2496–2509, 2020. 2
- [16] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 5, 7
- [17] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*, pages 35–51, 2018. 2
- [18] Xiao Liu, Pedro Sanchez, Spyridon Thermos, Alison Q O’Neil, and Sotirios A Tsaftaris. Learning disentangled representations in the imaging domain. *Medical Image Analysis*, page 102516, 2022. 2
- [19] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 6
- [20] Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. Smil: Multimodal learning with severely missing modality. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2302–2310, 2021. 1, 2, 5, 6, 7
- [21] 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
- [22] Miguel Monteiro, Loïc Le Folgoc, Daniel Coelho de Castro, Nick Pawlowski, Bernardo Marques, Konstantinos Kamnitsas, Mark van der Wilk, and Ben Glocker. Stochastic segmentation networks: Modelling spatially correlated aleatoric uncertainty. *Advances in Neural Information Processing Systems*, 33:12756–12767, 2020. 2
- [23] Mandela Patrick, Yuki M Asano, Polina Kuznetsova, Ruth Fong, Joao F Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations. *arXiv preprint arXiv:2003.04298*, 2020. 2
- [24] Mandela Patrick, Po-Yao Huang, Ishan Misra, Florian Metze, Andrea Vedaldi, Yuki M Asano, and João F Henriques. Space-time crop & attend: Improving cross-modal video representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10560–10572, 2021. 2
- [25] 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. 1, 2, 7

- [26] Gijs van Tulder and Marleen de Bruijne. Learning cross-modality representations from multi-modal images. *IEEE transactions on medical imaging*, 38(2):638–648, 2018. [2](#)
- [27] Valentin Vielzeuf, Alexis Lechervy, Stéphane Pateux, and Frédéric Jurie. Centralnet: a multilayer approach for multimodal fusion. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. [5](#)
- [28] Hu Wang, Qi Wu, and Chunhua Shen. Soft expert reward learning for vision-and-language navigation. In *European Conference on Computer Vision*, pages 126–141. Springer, 2020. [1](#)
- [29] Hu Wang, Jianpeng Zhang, Yuanhong Chen, Congbo Ma, Jodie Avery, Louise Hull, and Gustavo Carneiro. Uncertainty-aware multi-modal learning via cross-modal random network prediction. *arXiv preprint arXiv:2207.10851*, 2022. [1](#), [2](#)
- [30] Yikai Wang, Wenbing Huang, Fuchun Sun, Tingyang Xu, Yu Rong, and Junzhou Huang. Deep multimodal fusion by channel exchanging. *Advances in Neural Information Processing Systems*, 33:4835–4845, 2020. [2](#)
- [31] 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](#), [2](#), [7](#)
- [32] Zhangjing Wang, Yu Wu, and Qingqing Niu. Multi-sensor fusion in automated driving: A survey. *Ieee Access*, 8:2847–2868, 2019. [1](#)
- [33] Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton Van Den Hengel. Ask me anything: Free-form visual question answering based on knowledge from external sources. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4622–4630, 2016. [1](#)
- [34] Qiyue Yin, Shu Wu, and Liang Wang. Unified subspace learning for incomplete and unlabeled multi-view data. *Pattern Recognition*, 67:313–327, 2017. [2](#)
- [35] 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. *arXiv preprint arXiv:2206.02425*, 2022. [2](#), [6](#)