

# Unified Cross-Modal Medical Image Synthesis with Hierarchical Mixture of Product-of-Experts

Reuben Dorent, Nazim Haouchine, Alexandra Golby, Sarah Frisken, Tina Kapur, William Wells

**Abstract**—We propose a deep mixture of multimodal hierarchical variational auto-encoders called MMHVAE that synthesizes missing images from observed images in different modalities. MMHVAE’s design focuses on tackling four challenges: (i) creating a complex latent representation of multimodal data to generate high-resolution images; (ii) encouraging the variational distributions to estimate the missing information needed for cross-modal image synthesis; (iii) learning to fuse multimodal information in the context of missing data; (iv) leveraging dataset-level information to handle incomplete data sets at training time. Extensive experiments are performed on the challenging problem of pre-operative brain multi-parametric magnetic resonance and intra-operative ultrasound imaging.

**Index Terms**—Missing Data, Hierarchical Variational Autoencoders, Image Synthesis, Multimodal Image Computing

## 1 INTRODUCTION

MISSING data is a widespread problem in image analysis that presents significant practical and methodological challenges. This issue is common when the data collection is *longitudinal* and *multimodal*. For example, in medical image analysis, different sensors (e.g., MRI, ultrasound imaging) collect patient data at various stages (diagnosis, monitoring, surgery, follow-up), often resulting in incomplete image sets [1]. These challenges also arise in dynamic environments like remote sensing [2] or automated driving [3], where sensor availability and type may change over time. Consequently, handling missing data is essential for robust analysis in these multimodal scenarios.

The statistical literature offers an extensive set of data imputation methods to tackle missing data [4]. These methods range from simple techniques, such as zero or mean imputation, to more advanced ones, like multiple imputation [5] and K-nearest neighbors (KNN) [6]. However, these methods primarily address missing data in tabular datasets and are typically not applicable to complex, high-dimensional data such as images. This limitation is particularly pronounced in the context of deep learning models since these models exhibit a strong sensitivity to subtle changes in the data distribution [7]. For these reasons, addressing missing data in image-based datasets requires techniques designed to tackle the complexity of images.

To tackle missing imaging data, deep-learning-based methods have been proposed to perform cross-modal image synthesis [8], where missing data is estimated using observed data from other modalities. These methods include conditional generative adversarial networks (GANs) [9], [10], [11], [12] and conditional variational auto-encoders [13]. Yet, these synthesis frameworks do

not have the flexibility to handle missing data, requiring training for each possible combination of observed data.

To address these limitations, unified approaches have been proposed to perform inference from partially observed data. To have the flexibility for handling incomplete data, most approaches use zero imputation and are trained using artificially zero-imputed data [14], [15], [16], [17]. Other methods bypass imputation, focusing instead on creating a common feature space for images across different modalities. This involves extracting and fusing modality-specific features through arithmetic operations like mean [18], [19], [20], max [21], or a combination of sum, product and max [22]. However, these methods often fail to encourage the network to learn a shared multimodal latent representation and need more theoretical foundations. A more principled solution is offered by multimodal Variational Auto-Encoders (MVAEs) [23], [24], which utilize a probabilistic fusion operation to create a shared representation space. However, the low-dimensional nature of the latent space in MVAEs typically leads to the generation of blurry synthetic images. In contrast, Hierarchical VAEs (HVAEs) [25], [26], [27], [28] have demonstrated improved generative performance by leveraging a more complex shared latent representation. Yet, HVAEs are designed for mono-modal settings and do not handle missing data. Another key common limitation across all these methods is the requirement for complete data during training, which often results in excluding a large part of the available data.

In this work, we introduce a hierarchical multimodal variational auto-encoder (MMHVAE) that performs cross-modal image synthesis and is designed to handle incomplete data at both training and inference times. **The main contributions are summarized as follows:**

- 1) We propose a novel hierarchical multimodal variational auto-encoder to perform unified cross-modality image synthesis and allow for an incomplete set of images as input.
- 2) We model the variational posterior as a mixture

• R. Dorent, N. Haouchine, A. Golby, S. Frisken, T. Kapur and W. Wells are with Harvard Medical School and the Brigham and Women’s Hospital, MA, USA. R. Dorent is also with MIND Team, Inria Saclay, Université Paris-Saclay, Palaiseau, France and Sorbonne Université, Institut du Cerveau - Paris Brain Institute - ICM, CNRS, Inria, Inserm, AP-HP, Hôpital de la Pitié Salpêtrière, F-75013, Paris, France. W. Wells is also with Massachusetts Institute of Technology, MA, USA.

of Product-of-Experts. Each Product-of-Experts has a factorization similar to the true posterior. We demonstrate that this approach encourages the Product-of-Experts to not only encode observed information but also estimate missing information needed for image synthesis.

- 3) In contrast to most existing approaches, we assume that training data may not always be complete. **To regularize the distributions on the non-observed part of the training data, we introduce an adversarial strategy to leverage dataset-level information.**
- 4) Intensive experiments are performed on the challenging problem of cross-modality image synthesis between multi-parametric Magnetic Resonance sequences and ultrasound data. This includes an evaluation of image synthesis using paired multimodal data and two relevant downstream tasks: image segmentation and image registration. **Better image synthesis is obtained compared to existing unified cross-modality approaches while being less computationally heavy.**

**This work is a substantial extension of our conference paper [29].** Improvements include: 1) Generalization to more than two imaging modalities; 2) A novel method to leverage incomplete data at training time; 3) Additional mathematical proofs; 4) New experiments, including an ablation study and validation on two downstream tasks.

## 2 RELATED WORK

In this section, we review Variational Auto-Encoders (VAEs), their hierarchical and multimodal extensions and state-of-the-art unified cross-modal synthesis frameworks.

### 2.1 Variational Auto-Encoders (VAEs).

The goal of Variational Auto-Encoders (VAEs) [30] is to train a generative model in the form of  $p(\mathbf{x}, \mathbf{z}) = p(\mathbf{z})p(\mathbf{x}|\mathbf{z})$  where  $p(\mathbf{z})$  is a prior distribution (e.g. isotropic Normal distribution) over latent variables  $\mathbf{Z} \in \mathbb{R}^H$  and where  $p_\theta(\mathbf{x}|\mathbf{z})$  is a decoder parameterized by  $\theta$  that generates data  $\mathbf{X} \in \mathbb{R}^\Omega$  given  $\mathbf{Z}$ . Its graphical model is shown in Fig.1. The latent space dimension,  $H$ , is typically much lower than the image space dimension  $\Omega$ , i.e.  $H \ll \Omega$ . The training goal with respect to  $\theta$  is to maximize the marginal likelihood of the data  $p_\theta(\mathbf{x})$  (the ‘‘evidence’’); however since the true posterior  $p_\theta(\mathbf{z}|\mathbf{x})$  is in general intractable, the variational evidence lower bound (ELBO) is instead optimized. The ELBO  $\mathcal{L}_{\text{VAE}}(\mathbf{x}; \theta, \phi)$  is defined by introducing an approximate posterior  $q_\phi(\mathbf{z}|\mathbf{x})$  with parameters  $\phi$ :

$$\mathcal{L}_{\text{VAE}}(\mathbf{x}; \theta, \phi) \triangleq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \text{KL}[q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})], \quad (1)$$

where  $\text{KL}[q||p]$  is the Kullback-Leibler divergence between distributions  $q$  and  $p$ .

### 2.2 Hierarchical Variational Auto-Encoders (HVAEs).

To increase the expressiveness of both the approximate posterior and prior, Hierarchical VAEs (HVAEs)

have been introduced [25], [26], [27], [28]. In these extensions, the latent variable  $\mathbf{Z}$  is partitioned into disjoint groups, i.e.  $\mathbf{Z} = \{Z_1, \dots, Z_L\}$ , where  $L$  is the number of groups. The dimension of the latent representations  $Z_l$  typically exponentially decreases with the depth  $l$ . Then, the prior distribution is represented by  $p_\theta(\mathbf{z}) = p(z_L) \prod_{l=1}^{L-1} p_\theta(z_l|\mathbf{z}_{>l})$ , where  $\mathbf{z}_{>l} = (z_k)_{k=l+1}^L$ , and the approximate posterior by  $q_\phi(\mathbf{z}|\mathbf{x}) = q_\phi(z_L|\mathbf{x}) \prod_{l=1}^{L-1} q_\phi(z_l|\mathbf{z}_{>l}, \mathbf{x})$ . As HVAEs were originally designed to learn a mono-modal generative model, they assume that all the data  $\mathbf{X}$  is presented at all times. In contrast, our approach has the flexibility to handle incomplete sets of multimodal images at training and testing times, enabling practical cross-modal image synthesis.

### 2.3 Multimodal Variational Auto-Encoders (MVAEs)

Multimodal VAEs (MVAEs) [23], [24], [31] extended VAEs by learning a multimodal data representation and supporting missing data at inference time. MVAEs assume that  $M$  paired images  $\mathbf{X} = (X_1, \dots, X_M) \in \mathbb{R}^{M \times \Omega}$  are conditionally independent given a shared representation  $\mathbf{Z}$  as highlighted in Fig.1, i.e.  $p_\theta(\mathbf{x}|\mathbf{z}) = \prod_{i=1}^M p(x_i|\mathbf{z})$ .

Instead of training one single variational network  $q_\phi(\mathbf{z}|\mathbf{x})$  that requires all images to be present at all times, MVAEs factorize the approximate posterior as a combination of  $M$  unimodal variational distributions  $(q_\phi(\mathbf{z}|x_j))_{j=1}^M$ . Specifically, MVAE [23] factorizes the variational posterior as a Product-of-Experts (PoE), i.e.:

$$q_\phi^{\text{PoE}}(\mathbf{z}|\mathbf{x}) = p(\mathbf{z}) \prod_{j=1}^M q_\phi(\mathbf{z}|x_j). \quad (2)$$

Alternatively, MMVAE [31] proposed to decompose the variational posterior using a sum rule as a Mixture-of-Experts (MoE), i.e.:

$$q_\phi^{\text{MoE}}(\mathbf{z}|\mathbf{x}) = \frac{1}{M} \sum_{j=1}^M q_\phi(\mathbf{z}|x_j). \quad (3)$$

While the sum rule has been shown to be most resilient to estimation errors [32], the product rule follows a factorization similar to the true posterior [23].

Although the approximate posteriors of MVAEs have the flexibility to encode incomplete observations, these methods both face limitations in performing cross-modal image synthesis. The product rule typically generates misaligned latent representations across modalities, necessitating an ad hoc training sampling procedure to ensure similar representations for incomplete observations. In contrast, the sum rule focuses on generating all modalities from a single modality but does not combine representations when more than one modality is available. To address these limitations, MoPoE [33] introduced a variational posterior modeled as a mixture of product-of-experts. However, similar to MVAE and MMVAE, it relies on low-dimensional latent representations, which are inadequate for high-resolution image synthesis. Moreover, these methods assume access to all modalities during training, limiting their ability to leverage incomplete observed datasets. To overcome these limitations, we propose in this work a model that (i) learns




Fig. 1. Graphical models of: (a) variational auto-encoder (VAE); (b) hierarchical VAE (HVAE); (c) multimodal VAE (MVAE); (d) Our mixture of multimodal hierarchical VAE with missing data. Observed variables are in grey. In our model, variables are not always observed (partially gray).

to effectively combine representations from multiple input modalities, (ii) accurately estimates missing information for cross-modal image synthesis, (iii) employs a more complex latent representation to enable high-resolution image synthesis, and (iv) leverages incomplete training data to learn more robust and generalizable representations.

## 2.4 Unified cross-modal image synthesis

Several methodologies employing Generative Adversarial Networks (GANs) [14], [16], [17] and Transformers [15] have been proposed for unified cross-modal image synthesis, where a single model is utilized to synthesize all imaging modalities given any combination of incomplete observed data. However, these approaches omit the key aspect of learning a shared representation among multimodal data. Typically, these techniques concatenate multimodal images and substitute the missing images with zero tensors, i.e., zero-imputation. During training, all imaging data is assumed to be available. More specifically, the training procedure involves randomly substituting images with zero tensors and learning to synthesize the substituted images using supervised learning. More recently, cross-modal image synthesis frameworks based on diffusion models have been proposed. For instance, Cola-DIFF [34] employs a latent diffusion model specifically designed for brain MRI, but it only supports a many-to-one synthesis setting, making it unsuitable for unified synthesis. In contrast, some unified diffusion-based approaches have been introduced, where missing modalities are initialized with random noise and iteratively refined in the image space [35], [36], [37]. However, these methods do not learn compact latent representations and typically require large training datasets to prevent memorization [38]. They are also computationally intensive at inference time. Finally, like other unified approaches based on GANs or Transformers, they assume access to complete training data. In contrast to all these methods, our framework is designed to learn a common representation of multimodal data and has the flexibility to leverage incomplete training sets.

## 3 MIXTURE OF MULTIMODAL HIERARCHICAL VARIATIONAL AUTO-ENCODERS

In this paper, we propose a deep mixture of multimodal hierarchical VAE called MMHVAE that synthesizes missing modalities from observed images in different modalities. MMHVAE’s design focuses on tackling four challenges: (i) creating a complex latent representation of multimodal data to generate high-resolution images; (ii) encouraging the variational distributions to estimate the missing information needed for cross-modal image synthesis; (iii) learning to fuse multimodal information in the context of missing data; (iv) leveraging dataset-level information to handle incomplete data sets at training time.

### 3.1 Hierarchical latent representation of multimodal images

Let the random variable  $\mathbf{X} = (X_1, \dots, X_M) \in \mathbb{R}^{M \times \Omega}$  be a complete set of paired (i.e. co-registered) multimodal images, where  $M$  is the total number of image modalities and  $\Omega$  the number of pixels. The images  $\mathbf{X}$  are assumed to be conditionally independent given a latent random variable  $\mathbf{Z}$ . Then, the conditional distribution  $p_\theta(\mathbf{x}|\mathbf{z})$  parameterized by  $\theta$  can be written as:

$$p_\theta(\mathbf{x}|\mathbf{z}) = \prod_{j=1}^M p_\theta(x_j|\mathbf{z}). \quad (4)$$

Given that VAEs and MVAEs typically produce blurry images, we propose to use a hierarchical representation of the latent variable  $\mathbf{Z}$  to increase the expressiveness of the model as in HVAEs. Specifically, the latent variable  $Z$  is partitioned into disjoint groups, as shown in Fig.1 i.e.  $\mathbf{Z} = (Z_1, \dots, Z_L)$ , where  $L$  is the number of groups. The prior  $p(\mathbf{z})$  is then represented by:

$$p_\theta(\mathbf{z}) = p(z_L) \prod_{l=1}^{L-1} p_{\theta_l}(z_l|z_{>l}), \quad (5)$$

where  $z_{>l} = (z_k)_{k=l+1}^L$ ,  $p(z_L) = \mathcal{N}(z_L; \mathbf{0}_{H_L}, I_{H_L})$  is an isotropic Normal prior distribution and the conditional prior distributions  $p_{\theta_l}(z_l|z_{>l})$  are Normal distributions with

mean and diagonal covariance parameterized using neural networks, i.e.  $p_{\theta_l}(z_l|z_{>l}) = \mathcal{N}(z_l; \mu_{\theta_l}(z_{>l}), D_{\theta_l}(z_{>l}))$ . Note that the dimension of the latent representations exponentially decreases with the depth. In particular, the coarsest latent representation is a global descriptor of the multimodal images of size  $H_L$  (e.g.  $H_L = 128$ ), while the finest latent variable representation corresponds to a set of multimodal pixel descriptors of size  $F$ , i.e.  $H_1 = \Omega \times F$ . Thanks to the hierarchy, we can learn a complex prior for the high-dimensional representation  $z_1$  with complex local correlations characteristic of real images. In this work, we assume that the finest latent representation  $z_1$  carries all the information required to describe each image, i.e.:

$$p_{\theta}(\mathbf{x}|\mathbf{z}) = \prod_{j=1}^M p_{\theta_j}(x_j|z_1), \quad (6)$$

where the image decoding distributions are modeled as Normal with fixed variance  $\sigma$ , i.e.  $p_{\theta_j}(x_j|z_1) = \mathcal{N}(x_j; \mu_{\theta_j}(z_1), \sigma I)$ .

### 3.2 Marginal log-likelihood objective with incomplete data

In this work, we face the problem of missing data at inference and training time. We denote the vector of missingness indicators  $\mathbf{R}$  with  $R_j = 1$  if  $X_j$  is observed and 0 otherwise. The partially observed samples  $\mathbf{X}$  can be divided into an observed component  $\mathbf{X}_{\mathbf{R}}^o = \{X_i, \text{ such that } R_i = 1\}$  and a missing component  $\mathbf{X}_{\mathbf{R}}^m = \{X_i, \text{ such that } R_i = 0\}$ .

Let  $\mathcal{X} \times \mathcal{R} = \{(\mathbf{x}_1, \mathbf{r}_1), \dots, (\mathbf{x}_N, \mathbf{r}_N)\}$  be a training set of  $N$  partially observed images and  $\mathcal{X}^o = \{\mathbf{x}_{1;\mathbf{r}_1}^o, \dots, \mathbf{x}_{N;\mathbf{r}_N}^o\}$  its restriction to the observed component.

The overall objective is to maximize the expected observed marginal log-likelihood, i.e.  $\mathbb{E}_{(\mathbf{x}, \mathbf{r}) \sim p_{\text{data}}}[\log p_{\theta}(\mathbf{x}_{\mathbf{r}}^o)]$ . As the true posterior is intractable, the expected value of a tractable lower-bound is instead maximized. While the expected value is estimated using the samples in  $\mathcal{X}^o$ , a variational distribution  $q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}}^o)$  is introduced to approximate the posterior  $p_{\theta}(\mathbf{z}|\mathbf{x}_{\mathbf{r}}^o)$ .

### 3.3 Modeling the variational posterior as a mixture

In this work, we propose to model the variational posterior as a mixture distribution, where each component approximates the true posterior with incomplete input data. We show that this encourages the mixture components to encode all available information and estimate the missing information needed for cross-modal image synthesis.

Let  $S$  be the set of all  $2^M - 1$  vectors of length  $M$  consisting solely of zeros and ones, with each vector having at least one non-zero entry. For any vector  $\mathbf{r} \in S$ , let  $S_{\mathbf{r}}$  be the subset of indicators in  $S$  that share the zero indices of  $\mathbf{r}$ , i.e.  $S_{\mathbf{r}} = \{\mathbf{r}' \in S, \text{ s.t.: } r'_i = 0 \text{ if } r_i = 0\}$  (e.g.  $S_{[1,1,0,0]} = \{[1, 1, 0, 0], [1, 0, 0, 0], [0, 1, 0, 0]\}$ ).

We propose to express the variational  $q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}}^o)$  as a mixture (convex combination) of distributions represented by:

$$q_{\phi}^{\text{MMHVAE}}(\mathbf{z}|\mathbf{x}_{\mathbf{r}}^o) = \sum_{\mathbf{r}' \in S_{\mathbf{r}}} \alpha_{\mathbf{r}'}^{(\mathbf{r})} q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o) \quad (7)$$

where the input data  $\mathbf{x}_{\mathbf{r}'}^o$  is a subset of the observations  $\mathbf{x}_{\mathbf{r}}^o$  and  $\alpha_{\mathbf{r}'}^{(\mathbf{r})}$  are the mixture weights such that  $\alpha_{\mathbf{r}'}^{(\mathbf{r})} \geq 0$  and  $\sum_{\mathbf{r}' \in S_{\mathbf{r}}} \alpha_{\mathbf{r}'}^{(\mathbf{r})} = 1$ . These weights are hyper-parameters.

Let  $p_{\alpha}(\cdot|\mathbf{R} = \mathbf{r})$  denote the distribution of the mixture components  $\mathbf{R}'$  given  $\mathbf{R}$ , i.e.  $p_{\alpha}(\mathbf{r}'|\mathbf{r}) = \alpha_{\mathbf{r}'}^{(\mathbf{r})}$ . We can demonstrate that the evidence  $\log p(\mathbf{x}_{\mathbf{r}}^o)$  is lower-bounded by a new ELBO  $\mathcal{L}_{\text{MMHVAE}}^{\text{ELBO}}(\mathbf{x}_{\mathbf{r}}^o; \theta, \phi)$  defined as follows:

$$\log p(\mathbf{x}_{\mathbf{r}}^o) \geq \mathcal{L}_{\text{MMHVAE}}^{\text{ELBO}}(\mathbf{x}_{\mathbf{r}}^o; \theta, \phi) \triangleq \mathbb{E}_{\mathbf{r}' \sim p_{\alpha}(\cdot|\mathbf{r})} [\mathcal{L}(\mathbf{x}_{\mathbf{r}}^o, \mathbf{r}'; \theta, \phi)] \quad (8)$$

where:

$$\begin{aligned} \mathcal{L}(\mathbf{x}_{\mathbf{r}}^o, \mathbf{r}'; \theta, \phi) = & \sum_{\substack{j=1 \\ \text{s.t. } r'_j=1}}^M \underbrace{\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)} [\log p_{\theta}(x_j|\mathbf{z})]}_{\text{reconstruction input } x_j} \\ & + \sum_{\substack{j=1 \\ \text{s.t. } (1-r'_j)r_j=1}}^M \underbrace{\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)} [\log p_{\theta}(x_j|\mathbf{z})]}_{\text{synthesis of } x_j \text{ using } \mathbf{x}_{\mathbf{r}'}^o} \\ & - \underbrace{\text{KL}[q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o) \| p_{\theta}(\mathbf{z})]}_{\text{regularization}}. \end{aligned} \quad (9)$$

*Proof.* See Appendix 1.1. ■

The proposed ELBO  $\mathcal{L}_{\text{MMHVAE}}^{\text{ELBO}}(\mathbf{x}_{\mathbf{r}}^o; \theta, \phi)$  is the expected value of  $\mathcal{L}(\mathbf{x}_{\mathbf{r}}^o, \mathbf{R}'; \theta, \phi)$ , which contains input reconstruction, cross-modal image synthesis, and regularization terms.

This formulation not only allows optimization of a tractable upper-bound of the log likelihood, it also encourages the variational posterior to estimate all the information required to generate all the images. As demonstrated in Appendix 1.2, the ELBO  $\mathcal{L}_{\text{MMHVAE}}^{\text{ELBO}}(\mathbf{x}_{\mathbf{r}}^o; \theta, \phi)$  is in fact maximal with respect to the variational parameters  $\phi$  when the variational components  $q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)$  are equal to the true posterior distribution  $p_{\theta}(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)$ . Consequently, the objective of the variational distribution  $q_{\phi}(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)$  is to encode the information needed to generate the input data ( $\{x_i, \text{ s.t. } r'_i = 1\}$ ) while estimating the missing information to generate the non-encoded ones ( $\{x_i, \text{ s.t. } r'_i = 0 \text{ and } r_i = 1\}$ ). This especially encourages the latent representations to be aligned across any subset of modalities.

### 3.4 Variational parameterization for fusing incomplete multimodal inputs

Our next objective is to construct a variational component for each non-empty subset of observed data that approximates the posterior distribution. A straightforward approach would involve creating a variational distribution for every possible set of input data, resulting in the handling of  $2^M - 1$  encoding networks. Inspired by MVAEs [23], [31], we instead propose to create  $M$  unimodal encoding networks. This section demonstrates that we can identify candidates for the variational components that can be expressed (i) with a factorization similar to the true posterior  $p_{\theta}(\mathbf{z}|\mathbf{x}_{\mathbf{r}}^o)$ ; (ii) as a combination of variational unimodal distributions; and (iii) by fusing input information using a principled operation at each level of the hierarchy.




Fig. 2. The neural networks implementing (a) the encoder  $q(\mathbf{z}|\mathbf{x})$  and decoder  $p(\mathbf{x}|\mathbf{z}_1)$ ; (b) the encoder  $q(\mathbf{z}|x_1)$  and decoder  $p(x_2|\mathbf{z}_1)$  for a  $L = 3$  group hierarchical VAE with  $M = 2$  modalities.

Let  $\mathbf{x}_r^o$  be a set of observed images. The true hierarchical posterior  $p_\theta(\mathbf{z}|\mathbf{x}_r^o)$  can be factorized as a hierarchical combination of the conditional unimodal distributions  $((p_\theta(z_l|x_j, \mathbf{z}_{>l}))_{j=1 \text{ s.t. } r'_j=1}^M)_{l=1}^L$  and the prior distribution  $(p_{\theta_l}(z_l|\mathbf{z}_{>l}))_{l=1}^L$  as follows:

$$p_\theta(\mathbf{z}|\mathbf{x}_r^o) \propto \underbrace{\left( p(z_L) \prod_{j=1 \text{ s.t. } r'_j=1}^M \frac{\text{unimodal } p_\theta(z_L|x_j)}{p(z_L)} \right)}_{\propto p_\theta(z_L|\mathbf{x}_r^o)} \prod_{l=1}^{L-1} \underbrace{\left( p_{\theta_l}(z_l|\mathbf{z}_{>l}) \prod_{j=1 \text{ s.t. } r'_j=1}^M \frac{\text{unimodal } p_\theta(z_l|x_j, \mathbf{z}_{>l})}{p_{\theta_l}(z_l|\mathbf{z}_{>l})} \right)}_{\propto p_\theta(z_l|\mathbf{x}_r^o, \mathbf{z}_{>l})}. \quad (10)$$

*Proof.* See Appendix 1.3. ■

Unlike the conditional prior distributions  $p_{\theta_l}(z_l|\mathbf{z}_{>l})$  parameterized using neural networks in 3.1, the conditional unimodal distributions  $p_\theta(z_l|x_j, \mathbf{z}_{>l})$  are intractable. We thus propose a variational approach. Specifically, let  $\mathbf{r}'$  be an indicator vector in  $S_r$ . We propose to factorize the variational posterior distribution  $q(\mathbf{z}|\mathbf{x}_r^o)$  similarly to the true posterior  $p_\theta(\mathbf{z}|\mathbf{x}_r^o)$  defined

in (10) as a product of conditional prior distributions  $p_{\theta_l}(z_l|\mathbf{z}_{>l})$  and variational unimodal distributions  $(q(z_l|x_j, \mathbf{z}_{>l}))_{j \in \{1, \dots, M\} \text{ s.t. } r'_j=1}$ :

$$q(\mathbf{z}|\mathbf{x}_r^o) \propto \underbrace{\left( p(z_L) \prod_{j=1 \text{ s.t. } r'_j=1}^M \frac{q(z_L|x_j)}{p(z_L)} \right)}_{q(z_L|\mathbf{x}_r^o)} \prod_{l=1}^{L-1} \underbrace{\left( p_{\theta_l}(z_l|\mathbf{z}_{>l}) \prod_{j=1 \text{ s.t. } r'_j=1}^M \frac{q(z_l|x_j, \mathbf{z}_{>l})}{p_{\theta_l}(z_l|\mathbf{z}_{>l})} \right)}_{q_{\phi_l, \theta_l}(z_l|\mathbf{x}_r^o, \mathbf{z}_{>l})}. \quad (11)$$

While the product of two Normal densities  $p_1$  and  $p_2$  is proportional to a Normal density with closed form solutions for its parameters, the ratio  $\frac{p_1}{p_2}$  is a Normal density if and only if the  $p_1$ 's variance is element-wise larger than  $p_2$ 's variance. To satisfy this constraint, we propose a parametrization of the approximate conditional unimodal distributions  $q(z_l|x_j, \mathbf{z}_{>l})$  that guarantees  $\frac{q(z_l|x_j, \mathbf{z}_{>l})}{p_{\theta_l}(z_l|\mathbf{z}_{>l})}$  to always be a Normal distribution. Specifically, the approximate conditional unimodal distributions  $q(z_l|x_j, \mathbf{z}_{>l})$  are defined

as a product of the conditional prior  $p_{\theta_l}(z_l|z_{>l})$  and residual conditional unimodal Normal distributions  $h_{\phi_l^j}(z_l|x_j, z_{>l})$ :

$$q(z_l|x_j, z_{>l}) = p_{\theta_l}(z_l|z_{>l})h_{\phi_l^j}(z_l|x_j, z_{>l}), \quad (12)$$

where  $h_{\phi_l^j}(z_l|x_j, z_{>l})$  is a factorized Normal distribution parameterized using neural networks, i.e.  $h_{\phi_l^j}(z_l|x_j, z_{>l}) = \mathcal{N}(z_l; \mu_{\phi_l^j}(x_j, z_{>l}); D_{\phi_l^j}(x_j, z_{>l}))$ .

Consequently, the expression of the variational distribution  $q(z|\mathbf{x}_{r'}^o)$  in (11) can be simplified as:

$$q_{\phi, \theta}(z|\mathbf{x}_{r'}^o) \propto \underbrace{\left( p(z_L) \prod_{\substack{j=1 \\ \text{s.t. } r'_j=1}}^M h_{\phi_L}(z_L|x_j) \right)}_{q_{\phi_L, \theta_L}(z_L|\mathbf{x}_{r'}^o)} \prod_{l=1}^{L-1} \underbrace{\left( p_{\theta_l}(z_l|z_{>l}) \prod_{\substack{j=1 \\ \text{s.t. } r'_j=1}}^M h_{\phi_l^j}(z_l|x_j, z_{>l}) \right)}_{q_{\phi_l, \theta_l}(z_l|\mathbf{x}_{r'}^o, z_{>l})}. \quad (13)$$

The factorization in (13) shows that we can express all the  $2^M - 1$  variational posteriors using only  $M$  unimodal variational distributions. Importantly, a closed form solution allows merging unimodal contributions from each available modality at each level  $l \in \{1, \dots, L-1\}$  of the hierarchy. Specifically, the conditional variational distribution  $q_{\phi_l, \theta_l}(z_l|\mathbf{x}_{r'}^o, z_{>l})$  is a Normal distribution with mean  $\mu_{\phi_l, \theta_l}(\mathbf{x}_{r'}^o, z_{>l})$  and diagonal covariance  $D_{\phi_l, \theta_l}(\mathbf{x}_{r'}^o, z_{>l})$  defined by:

$$\begin{cases} D_{\phi_l, \theta_l}(\mathbf{x}_{r'}^o, z_{>l}) = \left( D_{\theta_l}(z_{>l})^{-1} + \sum_{\substack{j=1 \\ \text{s.t.} \\ r'_j=1}}^M D_{\phi_l^j}(x_j, z_{>l})^{-1} \right)^{-1} \\ \mu_{\phi_l, \theta_l}(\mathbf{x}_{r'}^o, z_{>l}) = D_{\phi_l, \theta_l}(\mathbf{x}_{r'}^o, z_{>l})^{-1} \left( D_{\theta_l}(z_{>l})^{-1} \mu_{\theta_l}(z_{>l}) + \sum_{\substack{j=1 \\ \text{s.t.} \\ r'_j=1}}^M D_{\phi_l^j}(x_j, z_{>l})^{-1} \mu_{\phi_l^j}(x_j, z_{>l}) \right) \end{cases} \quad (14)$$

with  $D_{\theta_L}(z_L) = I_{H_L}$  and  $\mu_{\theta_L}(z_L) = \mathbf{0}_{H_L}$

### 3.5 Regularizing image distributions using GAN loss

When maximizing the expected log-likelihood, our objective is to identify optimal parameter values that make the observed data most probable. However, the unobserved component of the data distribution  $p_{\theta}(\mathbf{X}_R^m)$  is disregarded, potentially resulting in out-of-distribution samples for the unobserved part of the data. To mitigate this issue, we propose to regularize the data distribution  $p_{\theta}(\mathbf{X})$  through adversarial learning. Specifically, for each modality  $j$ , a discriminator network  $d_{\psi_j}^j$ , parameterized by weights  $\psi_j$ , is trained to differentiate between real  $j$  images from  $\mathcal{X}^o$

**Algorithm 1** Training algorithm for MMHVAE.  $B$  is the batch size.  $T$  is the number of random subsets of observed images encoded at each training iteration.

---

```

1:  $\theta, \phi, \psi \leftarrow$  Initialize parameters
2: repeat
3:    $\mathbf{r} \leftarrow$  Random indicator (drawn from  $\mathcal{R}$  dataset)
4:    $\mathbf{x}_r^o \leftarrow$  Random minibatch of  $B$  multimodal images associated with indicator  $\mathbf{r}$  (drawn from  $\mathcal{X}$  dataset)
   Data model optimization
5:    $\mathcal{L}^{KL} \leftarrow 0$ 
6:    $\mathcal{L}^{Img} \leftarrow 0$ 
7:    $\mathcal{L}^{GAN} \leftarrow 0$ 
8:   for  $k \leftarrow 1$  to  $T$  do  $\triangleright$  Samples from the mixture
9:      $\mathbf{r}' \leftarrow$  Random indicator drawn from  $p_{\alpha}(\cdot|\mathbf{r})$ 
10:    for  $l \leftarrow L$  to  $1$  do  $\triangleright$  Encoding phase
11:       $z_l \leftarrow$  Random sample from  $q_{\phi_l, \theta_l}(z_l|\mathbf{x}_{r'}^o, z_{>l})$ 
12:       $\mathcal{L}_l^{KL} \leftarrow \text{KL}[q_{\phi_l, \theta_l}(z_l|\mathbf{x}_{r'}^o, z_{>l})\|p_{\theta_l}(z_l|z_{>l})]$ 
13:       $\mathcal{L}^{KL} \leftarrow \mathcal{L}^{KL} + \mathcal{L}_l^{KL}$ 
14:    end for
15:    for  $j \leftarrow 1$  to  $M$  do  $\triangleright$  Decoding phase
16:       $\hat{x}_j \leftarrow \mu_{\theta_j^x}(z_1)$ 
17:       $\mathcal{L}^{GAN} \leftarrow \mathcal{L}^{GAN} + d_{\psi_j}^j(\hat{x}_j)$ 
18:      if  $r_j = 1$  then
19:         $\mathcal{L}^{Img} \leftarrow \mathcal{L}^{Img} + \|x_j - \hat{x}_j\|^2$ 
20:      end if
21:    end for
22:  end for
23:   $\mathcal{L}^{Total} \leftarrow \mathcal{L}^{Img} + \lambda_{GAN}\mathcal{L}^{GAN} + \lambda_{KL}\mathcal{L}^{KL}$ 
24:   $\mathbf{g} \leftarrow \nabla_{\theta, \phi} \mathcal{L}^{Total}$ 
25:   $\theta, \phi \leftarrow$  Update parameters using gradient  $\mathbf{g}$ 
   Discriminators optimization
26:  for  $j \leftarrow 1$  to  $M$  do
27:    if  $r_j = 1$  then
28:       $m_j \leftarrow x_j$   $\triangleright$  Memory for future missing  $x_j$ 
29:    end if
30:     $\mathcal{L}^{Adv} \leftarrow d_{\psi_j}^j(m_j) - d_{\psi_j}^j(\hat{x}_j)$ 
31:     $\mathbf{g} \leftarrow \nabla_{\psi_j} \mathcal{L}^{Adv}$ 
32:     $\psi_j \leftarrow$  Update parameters using gradient  $\mathbf{g}$ 
33:  end for
34: until convergence of parameters  $(\theta, \phi, \psi)$ 

```

---

and  $j$  samples drawn from  $p_{\theta}$ , leading to the following minimization objective:

$$\min_{\psi_j} \mathbb{E}_{x_j} \left[ d_{\psi_j}^j(x_j) \right] - \mathbb{E}_{\mathbf{x}_r^o} \left[ \mathbb{E}_{q_{\phi}(z|\mathbf{x}_r^o)} \left[ d_{\psi_j}^j(\mu_{\theta_j^x}(z_1)) \right] \right] \quad (15)$$

Simultaneously, the data distribution model is trained to confound the discriminators  $(d_{\psi_j}^j)_{j=1}^M$ . Consequently, an additional GAN term  $\mathcal{L}^{GAN}(\mathbf{r}'; \theta, \phi)$  is introduced:

$$\mathcal{L}^{GAN}(\mathbf{r}'; \theta, \phi) = \sum_{j=1}^M \mathbb{E}_{q_{\phi}(z|\mathbf{x}_r^o)} \left[ d_{\psi_j}^j(\mu_{\theta_j^x}(z_1)) \right] \quad (16)$$

In particular, this GAN term encourages samples from the non-observed data distribution to be realistic by leveraging modality-specific knowledge at the dataset level.

### 3.6 Training strategy

Combining sections 3.2 and 3.5, the complete training objective is to maximize the expected log-likelihood using the expected lower-bound (9) while minimizing the expected regularization term (16), leading to the following minimization objective:

$$\min_{\theta, \phi} \mathbb{E}_{\mathbf{r} \sim p_{data}} \mathbb{E}_{\substack{\mathbf{x}_r \sim p_{data}(\cdot | \mathbf{r}) \\ \mathbf{r}' \sim p_{\alpha}(\cdot | \mathbf{r})}} \left[ \mathcal{L}^{Total}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi) \right] \quad (17)$$

where  $\mathcal{L}^{Total}$  is the sum of the negative ELBO defined in (9) and the regularization term in (16):

$$\mathcal{L}^{Total}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi) = -\mathcal{L}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi) + \lambda_{GAN} \mathcal{L}^{GAN}(\mathbf{r}'; \theta, \phi) \quad (18)$$

We use a stochastic gradient descent approach to minimize the expected loss  $\mathcal{L}^{Total}$ . First, we draw a random indicator  $\mathbf{r}$  from  $\mathcal{R}$ . Second, we draw a random mini-batch of size  $B$  of observed images from  $\mathcal{X}^o$  associated with the indicator  $\mathbf{r}$ . Third, we draw  $T$  vectors  $\mathbf{r}' \in S_r$  from  $p_{\alpha}(\cdot | \mathbf{r})$ . Finally, we minimize the average loss of  $\mathcal{L}^{Total}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi)$  for all drawn multimodal images  $\mathbf{x}_r^o$  in the mini-batch and all drawn vectors  $\mathbf{r}'$ . Note that the expected values in  $\mathcal{L}^{Total}$  are estimated using a unique sample at each level of the hierarchy as in [27]. Consequently, given a triplet  $(\mathbf{r}, \mathbf{x}_r^o, \mathbf{r}')$ , the total loss  $\mathcal{L}^{Total}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi)$  is defined as:

$$\begin{aligned} \mathcal{L}^{Total}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi) = & \sum_{j=1}^M \left\| \mu_{\theta_j^x}(z_1) - x_j \right\|^2 \\ & + \lambda_{KL} \sum_{l=1}^L \log \frac{q_{\phi_l, \theta_l}(z_l | \mathbf{x}_r^o, \mathbf{z}_{>l})}{p_{\theta_l}(z_l | \mathbf{z}_{>l})} \\ & + \lambda_{GAN} \sum_{j=1}^M d_{\psi_j}^j \left( \mu_{\theta_j^x}(z_1) \right), \end{aligned} \quad (19)$$

where  $\lambda_{KL}$  and  $\lambda_{GAN}$  weight the contributions of the KL divergence and GAN loss. The complete training strategy is presented in Algorithm 1.

## 4 EXPERIMENTAL SETUP

To assess the performance of our model, we conducted comprehensive experiments on the challenging task of cross-modal synthesis of multi-parametric magnetic resonance images (MRI) and intraoperative ultrasound images (iUS) in patients with brain tumors. Specifically, our focus was on the key MR sequences for brain tumor surgery, namely, the contrast-enhanced T<sub>1</sub>-weighted (ceT<sub>1</sub>), T<sub>2</sub>-weighted, and T<sub>2</sub> Fluid Attenuation Inversion Recovery (FLAIR) sequences, resulting in  $M = 4$  modalities. MMHVAE is empirically evaluated by (i) cross-modal image synthesis results with standard paired evaluation metrics, (ii) an ablation study analyzing each novel component, (iii) segmentation results as a downstream task, and (iv) image registration results as a downstream task.

### 4.1 Datasets

Three datasets were utilized in our experiments:

**1. ReMIND:** The Brain Resection Multimodal Imaging Database (ReMIND) comprises pre-operative multi-parametric and intra-operative ultrasound data collected

from 114 consecutive patients [1]. For this study, we selected all the patients ( $N = 104$ ) who underwent acquisition of both pre-operative 3D MRI and intraoperative pre-dural opening ultrasound (iUS), reconstructed from a tracked handheld 2D probe. Not all 3D MR sequences were acquired for each patient, resulting in 98 ceT<sub>1</sub>, 66 T<sub>2</sub>, and 21 FLAIR MR scans. In particular, only 13 patients have a complete set of MRI, highlighting the need for methods designed to handle missing data. To create a paired dataset, pre-operative MR images were affinely registered with the pre-dura iUS using NiftyReg [39], following the pipeline described in [40]. Three neurological experts manually checked the registration outputs. Images were resampled to an isotropic 0.5mm resolution and padded for an in-plane matrix of (192, 192). While the ultrasound images were acquired with the same ultrasound probe, preoperative MR images were collected under various clinical protocols and scanners at multiple institutions. This led to different intensity distributions for a given MR sequence.

**2. UPenn-GBM:** The UPenn-GBM dataset comprises multi-parametric preoperative MRI data for brain tumor patients with delineations of the distinct tumor sub-regions [41]. Ultrasound images are not available in this dataset. In contrast, all MR sequences were acquired pre-operatively for  $N = 611$  patients. All images are co-registered to the same anatomical template and resampled to an isotropic 1mm resolution. The brain images in UPenn-GBM are also included in the BraTS dataset [42], albeit without undergoing skull-stripping preprocessing. MR scans were acquired at a unique institution with different scanners and clinical protocols, leading to different intensity distributions for a given MR sequence.

**3. RESECT-SEG:** RESECT-SEG is a small dataset of 23 3D brain iUS with brain tumor delineations [43]. All images were acquired at a single institution with the same ultrasound probe. Note that this ultrasound probe differs from the one used in ReMIND, leading to a distribution gap between the iUS in RESECT-SEG and ReMIND data.

### 4.2 Implementation details

**Network Architecture:** Given that raw brain ultrasound images are typically 2D, we adopted a 2D U-Net-based architecture. The spatial resolution and feature dimension of the coarsest latent variable ( $z_L$ ) were chosen as  $1 \times 1$  and 256, respectively. The spatial and feature dimensions were successively doubled and halved after each level, resulting in a feature representation of dimension 8 for each pixel at group 1, denoted as  $z_1 \in \mathbb{R}^{192 \times 192 \times 8}$ . This results in  $L = 7$  latent variable levels. Following state-of-the-art architectures [25], residual cells from MobileNetV2 [44] are employed for the encoder and decoder, with Squeeze and Excitation [45] and Swish activation. The image decoders  $(\mu_{\theta_j^x})_{j=1}^M$  correspond to 5 ResNet blocks. At inference, we lower the temperature of the parametric distributions to 0.5, as performed in other HVAEs [25]. The code is available at <https://github.com/ReubenDo/MMHVAE>.

**Training Parameters:** The models are trained for 1000 epochs with a batch size of  $B = 16$ . The parameters  $\alpha_r^{(r)}$  are set such that the probability of drawing the full set of observable images (i.e.,  $\mathbf{r}' = \mathbf{r}$ ), the ultrasound image only,

or a set of MRI is equal (i.e.,  $p = \frac{1}{3}$ ). Additionally, a uniform distribution is chosen over the set of subsets of observable MRI. At each training iteration,  $T = 3$  indicator vectors  $\mathbf{r}'$  are drawn. The KL divergence is set to  $\lambda_{KL} = 0.001$ . The weight of the GAN loss is set to  $\lambda_{GAN} = 0$  for the 800 first epochs and then to  $\lambda_{GAN} = 0.025$  for the last 200 epochs. Models were trained on an A100 40GB GPU.

### 4.3 Competing methods

We compared our approach to three state-of-the-art approaches for unified cross-modal image synthesis. Curriculum learning from our model was used to ensure standard sample selection for all competing methods.

**MoPoE** [33]: We implemented the Mixture-of-Products-of-Expert (MoPoE) using our network architecture for fair comparison. The unification strategy is the same as Ours with a single latent variable ( $L = 1$ ).

**MM-GAN** [14]: MM-GAN is a unified synthesis model based on a convolutional GAN was considered. MM-GAN comprises CNN-based generator and discriminator networks, where the generator is based on U-Net. MM-GAN trains a single network under various source-target modality configurations. The original MM-GAN architecture was directly adopted. The unification strategy consists of concatenating all images and replacing missing ones with a zero tensor.

**ResViT** [15]: The current state-of-the-art unified synthesis model using transformers. The original ResViT architecture was directly adopted. The unification strategy in ResViT uses zero-imputation as MM-GAN.

**M2DN** [36]: The current state-of-the-art unified synthesis model based on diffusion models. M2DN applies a diffusion model with binary conditional codes to specify different imputation tasks. The unification strategy consists of concatenating all images and replacing missing ones with random noise. During the reverse diffusion step, the modalities are progressively reconstructed. To additionally perform harmonization, we adapted the original method using the Brownian bridge diffusion framework [46].

**MHVAE** [29]: MHVAE is our previously proposed model, which can handle only two modalities (iUS and  $T_2$ ).

## 4.4 Experiments design

To evaluate the effectiveness of our framework and compare it with state-of-the-art unified approaches, we conducted experiments on three tasks: (i) cross-modal image synthesis using paired data, (ii) downstream tumor segmentation, and (iii) downstream image registration. We also performed several ablation studies.

### 4.4.1 Harmonized cross-modal image synthesis

This task aims to synthesize harmonized images from all  $M = 4$  modalities given any incomplete set of paired non-harmonized images. Specifically, if the target image is available as input, the task is to perform image harmonization. If the target image is missing, the task is to synthesize the harmonized version of the missing image.

Medical data often acquired significant intensity distribution shifts when acquired at different centers with varying imaging protocols (e.g., different scanners and

acquisition parameters). For example, these intra-modality shifts can be observed for the MR data from ReMIND, as shown in Appendix Fig. 1. Through preliminary experiments, we found that performing cross-modality synthesis without considering these shifts leads to the synthesis of blurry MR images, mainly due to the intrinsic one-to-many contrast mapping between input and target images. To address this issue, data harmonization is required.

For data harmonization, we applied contrast linear normalization. We first shifted the minimal intensity value to 0. Then, we aligned the median intensity values within white matter, estimated using SynthSeg [47], to  $\frac{1}{7}$ ,  $\frac{1}{5}$  and  $\frac{1}{3}$  for  $T_2$ ,  $ceT_1$  and FLAIR respectively. As shown in Figure 6, this method aligns well the intensity distribution when median values are obtained from the 3D volumes. However, significant shifts were observed when normalization was performed slice by slice. Since our model takes 2D slices as input, we applied 3D contrast normalization only to the ground truth target images during training and testing, but not to the 2D input slices. Therefore, when a 2D MR slice is available as input, the task is to harmonize it.

Ultrasound images, on the other hand, depend on the probe’s position and angle, as they are generated from the reflection, scatter, and absorption of sound waves through tissues. To account for this spatial dependency during synthesis, all MR images were cropped to the field of view of the iUS probe, allowing us to encode the probe’s position and angle relative to the brain surface.

Since paired data is available for evaluation (ReMIND), we employed standard supervised evaluation metrics: PSNR (Peak Signal-to-Noise Ratio), SSIM (Structural Similarity), and LPIPS [48] (Learned Perceptual Image Similarity). We conducted a 3-fold cross-validation using stratified sampling to ensure that all possible combinations of inputs are adequately represented in each test set.

### 4.4.2 Ablation study

We performed an ablation study to assess the impact of three key components: (i) the hierarchical latent representation, (ii) the probabilistic fusion strategy for missing images, and (iii) the GAN-based regularization. These experiments were conducted on one fold.

First, we assessed the utility of the hierarchy in the latent representation. We compared our approach with  $L = 7$  latent variables to a non-hierarchical approach  $L = 1$  (group 7). Moreover, we compared our approach to three hierarchical approaches with  $L = 2$  (groups 1 and 7),  $L = 3$  (groups 1, 4, and 7), and  $L = 5$  (groups 1, 3, 4, 5, and 7) latent variables. A quantitative assessment was performed for the cross-modal image synthesis task presented above.

Second, we compared our proposed fusion operation with two common techniques for handling missing data.

**Concat. w/ zeros** uses the concatenation operation to merge information from multiple images. Similarly to MM-GAN [14] and ResViT [15], zero imputation is used when images are missing. In this model, only one encoder is used.

**Average** is a deterministic variant of our framework using the same network architecture. Features are extracted independently from each available image and averaged at each level of the hierarchy, as in [18], [19], [20].

TABLE 1

**Comparison against the state-of-the-art unified models for image synthesis.** Mean and standard deviation values are presented. For a given input, values in bold indicate the methods that achieved the best performance and were found to be statistically significantly better than all others, based on a paired Wilcoxon signed-rank test with Bonferroni correction ( $p < 0.01$ ). Arrows indicate the favorable direction of each metric.

	Input				iUS			T <sub>2</sub>			ceT <sub>1</sub>			FLAIR		
	iUS	T <sub>2</sub>	T <sub>1</sub>	FL	PSNR(dB)↑	SSIM(%)↑	LPIPS(%)↓	PSNR(dB)↑	SSIM(%)↑	LPIPS(%)↓	PSNR(dB)↑	SSIM(%)↑	LPIPS(%)↓	PSNR(dB)↑	SSIM(%)↑	LPIPS(%)↓
MoPoE [33]	•	○	○	○	22.4 (3.6)	75.7 (11.7)	31.6 (13.1)	<b>21.8 (5.0)</b>	<b>83.0 (8.6)</b>	24.1 (10.5)	<b>23.6 (5.3)</b>	<b>84.9 (7.6)</b>	20.2 (9.5)	18.4 (8.2)	<b>79.3 (10.3)</b>	22.9 (11.4)
MHVAE [29]	•	○	○	○	31.6 (2.2)	92.4 (3.3)	6.8 (2.6)	21.6 (4.0)	80.7 (9.1)	20.8 (8.8)	-	-	-	-	-	-
MM-GAN [14]	•	○	○	○	25.3 (5.0)	87.3 (7.3)	10.8 (5.6)	19.5 (4.6)	76.7 (11.0)	19.9 (8.2)	21.3 (4.6)	79.5 (9.4)	17.0 (7.8)	16.7 (7.7)	74.3 (11.9)	19.3 (9.7)
ResViT [15]	•	○	○	○	32.5 (4.1)	95.6 (2.6)	3.2 (1.7)	20.7 (4.2)	79.1 (9.8)	<b>16.0 (6.7)</b>	22.5 (5.0)	81.7 (8.4)	<b>13.5 (6.6)</b>	17.8 (7.5)	76.5 (10.8)	<b>14.5 (9.3)</b>
M2DN [36]	•	○	○	○	27.7 (5.5)	88.0 (7.1)	5.5 (2.8)	19.4 (5.1)	80.1 (9.5)	20.1 (8.9)	20.7 (5.7)	83.0 (8.1)	17.1 (8.3)	18.3 (7.5)	78.6 (10.2)	16.9 (10.0)
Ours	•	○	○	○	<b>46.3 (2.9)</b>	<b>99.6 (0.5)</b>	<b>0.2 (0.4)</b>	20.8 (4.3)	79.9 (9.6)	16.3 (6.6)	22.2 (5.4)	82.0 (8.5)	<b>14.0 (6.8)</b>	<b>18.6 (8.0)</b>	<b>79.0 (10.2)</b>	17.2 (9.5)
MoPoE [33]	○	•	○	○	<b>21.0 (4.4)</b>	<b>74.4 (12.4)</b>	31.6 (13.7)	21.8 (5.1)	83.2 (8.6)	24.0 (10.5)	24.4 (4.5)	85.9 (6.9)	19.3 (9.2)	20.3 (4.8)	81.0 (9.2)	21.4 (9.9)
MHVAE [29]	○	•	○	○	20.8 (3.7)	73.9 (12.0)	20.5 (8.5)	28.6 (3.1)	91.0 (3.9)	5.7 (2.1)	-	-	-	-	-	-
MM-GAN [14]	○	•	○	○	20.5 (4.4)	73.3 (12.5)	25.4 (10.6)	23.3 (4.4)	93.0 (3.4)	5.4 (2.1)	24.4 (3.5)	86.1 (6.4)	10.6 (4.9)	19.8 (3.8)	79.3 (9.7)	12.5 (5.4)
ResViT [15]	○	•	○	○	20.0 (4.2)	70.9 (13.3)	19.3 (7.9)	28.6 (4.6)	94.6 (3.0)	5.0 (2.7)	<b>24.9 (4.1)</b>	<b>86.9 (6.4)</b>	<b>8.8 (4.5)</b>	22.4 (3.4)	<b>83.5 (7.7)</b>	<b>8.5 (3.8)</b>
M2DN [36]	○	•	○	○	20.6 (4.9)	<b>74.5 (12.4)</b>	23.2 (9.3)	24.9 (5.3)	95.0 (3.1)	4.6 (2.5)	20.6 (4.8)	84.9 (7.3)	15.4 (7.5)	21.3 (4.4)	82.1 (8.6)	12.9 (6.1)
Ours	○	•	○	○	20.6 (4.2)	73.1 (12.4)	<b>18.9 (7.6)</b>	<b>35.9 (4.1)</b>	<b>98.5 (1.2)</b>	<b>1.1 (0.7)</b>	<b>24.9 (4.3)</b>	<b>86.7 (6.5)</b>	9.0 (4.4)	<b>23.5 (4.1)</b>	<b>85.1 (7.3)</b>	9.0 (4.4)
MoPoE [33]	○	○	•	○	20.9 (4.5)	73.8 (12.4)	32.3 (13.6)	21.7 (5.0)	82.8 (8.6)	24.3 (10.5)	23.6 (5.1)	85.0 (7.5)	20.4 (9.5)	18.3 (8.2)	79.2 (10.0)	23.9 (11.0)
MM-GAN [14]	○	○	•	○	20.0 (4.2)	72.4 (12.6)	26.3 (10.6)	21.2 (4.5)	82.1 (8.5)	13.6 (5.5)	28.1 (4.0)	94.8 (3.0)	4.4 (2.2)	18.6 (6.7)	78.1 (10.3)	14.3 (7.6)
ResViT [15]	○	○	•	○	19.8 (4.2)	70.0 (13.5)	<b>19.8 (8.0)</b>	22.6 (4.4)	82.6 (8.2)	12.1 (5.2)	29.6 (5.4)	94.7 (3.1)	4.3 (3.5)	20.1 (7.9)	81.6 (8.5)	<b>10.6 (8.5)</b>
M2DN [36]	○	○	•	○	19.8 (5.0)	73.1 (12.9)	24.8 (9.7)	19.3 (5.3)	80.4 (9.4)	20.3 (9.0)	24.7 (7.5)	93.5 (3.8)	6.7 (3.9)	18.5 (8.0)	78.9 (10.0)	17.4 (10.2)
Ours	○	○	•	○	<b>21.1 (4.2)</b>	<b>74.0 (11.9)</b>	20.9 (8.8)	<b>23.2 (4.4)</b>	<b>84.2 (7.9)</b>	<b>10.5 (4.6)</b>	<b>36.1 (6.2)</b>	<b>98.6 (2.0)</b>	<b>1.3 (3.0)</b>	<b>21.4 (7.8)</b>	<b>83.6 (7.9)</b>	<b>10.5 (7.3)</b>
MoPoE [33]	○	○	○	•	<b>20.9 (4.5)</b>	<b>74.8 (12.6)</b>	31.2 (13.8)	22.0 (4.9)	<b>84.1 (8.0)</b>	23.3 (10.1)	22.5 (7.2)	84.9 (7.8)	20.8 (10.2)	18.3 (7.8)	79.5 (9.7)	23.9 (11.5)
MM-GAN [14]	○	○	○	•	20.5 (4.3)	73.5 (12.8)	25.4 (11.0)	20.5 (4.3)	80.3 (9.4)	15.9 (6.9)	22.4 (7.0)	83.1 (8.4)	13.4 (7.9)	22.2 (7.2)	86.9 (7.3)	9.3 (6.9)
ResViT [15]	○	○	○	•	20.0 (4.3)	71.3 (13.5)	<b>19.2 (8.0)</b>	<b>22.6 (4.2)</b>	<b>83.4 (8.0)</b>	<b>11.5 (5.0)</b>	<b>22.7 (7.9)</b>	<b>85.2 (7.2)</b>	<b>10.6 (6.8)</b>	27.2 (8.6)	95.0 (5.2)	4.2 (8.3)
M2DN [36]	○	○	○	•	20.6 (5.0)	<b>74.8 (12.7)</b>	22.8 (9.3)	18.6 (5.2)	81.2 (9.0)	18.4 (8.1)	18.5 (8.2)	83.3 (8.2)	16.8 (8.4)	22.0 (8.7)	94.5 (4.7)	4.5 (6.3)
Ours	○	○	○	•	<b>20.3 (4.4)</b>	<b>73.4 (12.9)</b>	20.7 (9.1)	21.7 (4.4)	83.6 (8.1)	11.9 (4.8)	22.4 (7.5)	84.1 (7.9)	10.9 (7.4)	<b>31.3 (8.6)</b>	<b>97.6 (5.1)</b>	<b>2.5 (7.6)</b>
MoPoE [33]	○	•	•	○	<b>21.0 (4.5)</b>	<b>74.3 (12.4)</b>	31.9 (13.6)	21.8 (5.0)	82.9 (8.6)	24.3 (10.5)	24.4 (4.4)	85.9 (6.9)	19.3 (9.2)	20.2 (4.9)	80.3 (9.3)	22.1 (9.9)
MM-GAN [14]	○	•	•	○	20.7 (4.4)	73.6 (12.4)	24.9 (10.4)	24.0 (4.4)	93.3 (3.3)	5.1 (2.1)	28.2 (3.1)	93.9 (2.7)	5.0 (2.0)	20.4 (3.8)	80.6 (9.3)	12.0 (5.3)
ResViT [15]	○	•	•	○	20.1 (4.3)	71.1 (13.2)	19.0 (7.8)	28.2 (4.7)	93.9 (3.3)	5.2 (2.9)	29.2 (3.9)	93.6 (3.2)	4.6 (2.6)	22.8 (3.4)	84.6 (7.2)	<b>7.6 (3.5)</b>
M2DN [36]	○	•	•	○	20.6 (4.8)	<b>74.4 (12.2)</b>	23.0 (9.2)	24.9 (5.4)	95.0 (3.0)	4.5 (2.3)	24.9 (7.2)	93.9 (3.7)	5.4 (2.7)	21.6 (4.3)	82.7 (8.4)	12.1 (5.7)
Ours	○	•	•	○	<b>20.9 (4.0)</b>	73.6 (12.0)	<b>18.4 (7.4)</b>	<b>35.6 (4.5)</b>	<b>98.4 (1.3)</b>	<b>1.1 (0.7)</b>	<b>36.1 (4.6)</b>	<b>98.7 (0.7)</b>	<b>1.1 (0.6)</b>	<b>23.9 (3.8)</b>	<b>85.8 (6.9)</b>	8.6 (4.2)
MoPoE [33]	○	•	○	○	<b>21.3 (4.6)</b>	75.5 (12.5)	30.9 (14.0)	22.3 (4.9)	84.3 (8.0)	23.3 (10.1)	24.2 (4.4)	86.3 (6.6)	18.5 (8.5)	20.2 (4.7)	80.9 (9.2)	21.4 (9.8)
MM-GAN [14]	○	•	○	○	21.1 (4.5)	74.9 (12.3)	24.4 (10.9)	24.1 (4.2)	93.3 (3.3)	4.7 (2.0)	24.9 (3.5)	87.1 (6.1)	9.3 (4.3)	25.7 (4.2)	91.2 (4.6)	5.5 (2.5)
ResViT [15]	○	•	○	○	20.3 (4.3)	72.4 (13.2)	18.5 (7.9)	28.3 (4.5)	94.0 (3.2)	5.3 (3.0)	<b>25.1 (3.9)</b>	<b>88.2 (5.9)</b>	<b>7.8 (3.8)</b>	28.3 (3.6)	95.2 (2.4)	3.2 (2.1)
M2DN [36]	○	•	○	○	<b>21.2 (5.0)</b>	<b>75.8 (12.4)</b>	21.7 (9.1)	25.5 (5.1)	<b>95.5 (2.9)</b>	3.8 (2.0)	19.8 (5.3)	85.8 (7.0)	13.6 (6.1)	22.6 (6.9)	95.2 (2.6)	3.4 (1.9)
Ours	○	•	○	○	20.4 (4.0)	73.7 (12.5)	<b>18.1 (7.6)</b>	<b>35.1 (4.2)</b>	<b>98.3 (1.5)</b>	<b>1.4 (0.9)</b>	24.6 (3.9)	87.0 (6.1)	8.8 (4.0)	<b>33.3 (4.0)</b>	<b>98.2 (1.1)</b>	<b>1.3 (0.8)</b>
MoPoE [33]	○	○	•	•	<b>20.9 (4.6)</b>	<b>74.4 (12.6)</b>	31.8 (13.9)	22.0 (4.7)	83.6 (8.1)	24.1 (10.1)	22.8 (7.0)	85.2 (7.8)	20.3 (8.8)	18.2 (7.9)	79.1 (9.7)	24.3 (10.4)
MM-GAN [14]	○	○	•	•	20.5 (4.2)	73.4 (12.6)	25.3 (10.7)	21.4 (4.3)	82.2 (8.2)	14.9 (6.3)	26.9 (6.8)	93.2 (5.8)	5.3 (5.8)	23.0 (7.5)	89.0 (6.7)	7.7 (6.9)
ResViT [15]	○	○	•	•	20.0 (4.3)	71.2 (13.3)	<b>19.1 (7.9)</b>	23.4 (4.4)	84.7 (7.3)	10.4 (4.7)	27.2 (8.6)	93.1 (4.5)	5.6 (6.3)	26.3 (9.0)	94.2 (5.5)	4.7 (8.8)
M2DN [36]	○	○	•	•	20.6 (4.9)	<b>74.4 (12.5)</b>	22.9 (9.2)	18.5 (5.2)	81.7 (8.7)	17.1 (7.4)	22.7 (9.4)	93.7 (4.8)	5.7 (6.0)	22.4 (9.0)	94.6 (4.9)	4.5 (6.8)
Ours	○	○	•	•	<b>20.9 (4.0)</b>	74.0 (12.0)	19.4 (8.3)	<b>23.9 (4.0)</b>	<b>85.6 (6.9)</b>	<b>9.5 (3.8)</b>	<b>33.5 (9.8)</b>	<b>97.9 (4.0)</b>	<b>2.1 (6.2)</b>	<b>31.1 (8.9)</b>	<b>97.6 (4.0)</b>	<b>2.4 (6.7)</b>
MoPoE [33]	○	•	•	•	<b>21.2 (4.6)</b>	<b>75.0 (12.4)</b>	31.5 (13.9)	22.1 (4.9)	83.6 (8.1)	24.0 (10.1)	24.3 (4.3)	86.4 (6.6)	18.5 (8.4)	20.2 (4.7)	80.3 (9.2)	22.0 (9.8)
MM-GAN [14]	○	•	•	•	<b>21.1 (4.4)</b>	74.6 (12.2)	24.4 (10.7)	25.3 (4.0)	93.8 (3.0)	4.5 (1.9)	28.1 (3.5)	93.6 (3.0)	4.8 (2.0)	26.2 (3.9)	92.2 (4.1)	5.2 (2.4)
ResViT [15]	○	•	•	•	20.3 (4.3)	72.1 (13.0)	18.6 (7.8)	27.9 (4.7)	93.5 (3.4)	5.6 (3.2)	28.2 (3.9)	93.0 (3.5)	4.8 (2.4)	27.8 (3.5)	94.7 (2.7)	3.5 (2.2)
M2DN [36]	○	•	•	•	21.0 (4.9)	<b>75.1 (12.3)</b>	22.0 (9.0)	25.9 (5.0)	95.5 (2.5)	3.9 (1.9)	23.6 (6.9)	94.3 (3.5)	4.7 (2.4)	23.2 (6.7)	95.4 (2.3)	3.4 (1.8)
Ours	○	•	•	•	20.6 (3.9)	73.9 (12.1)	<b>18.0 (7.4)</b>	<b>35.1 (4.3)</b>	<b>98.2 (1.6)</b>	<b>1.3 (0.8)</b>	<b>34.2 (5.0)</b>	<b>98.4 (1.0)</b>	<b>1.2 (0.7)</b>	<b>32.8 (4.3)</b>	<b>98.1 (1.1)</b>	<b>1.5 (0.8)</b>
Average	MoPoE [33]	21.3 (4.3)	74.6 (12.3)	31.7 (13.5)	21.8 (5.0)	83.2 (8.5)	24.1 (10.4)	23.8 (5.2)	85.4 (7.3)	19.8 (9.3)	19.2 (6.8)	79.9 (9.6)	22.8 (10.6)			
	MM-GAN [14]	21.6 (4.9)	76.7 (13.0)	21.9 (11.5)	22.2 (4.8)	86.5 (10.0)	10.8 (7.9)	25.4 (5.2)	88.5 (8.8)	9.3 (7.2)	21.4 (6.8)	83.7 (10.6)	11.0 (7.9)			
	ResViT [15]	23.0 (6.8)	76.9 (15.8)	15.4 (9.8)	25.1 (5.5)	87.8 (9.2)	9.3 (6.3)	26.3 (5.9)	89.1 (7.8)	7.9 (6.1)	24.0 (7.5)	88.0 (10.0)	7.2 (7.8)			
	M2DN [36]	22.2 (6.0)	77.5 (12.9)	19.1 (11.3)	22.1 (6.0)	87.8 (10.1)	12.0 (10.0)	22.4 (7.0)	88.8 (7.9)	11.1 (8.1)	21.1 (7.5)	87.7 (10.3)	9.5 (9.0)			
	Ours	<b>27.0 (11.6)</b>	<b>79.9 (15.4)</b>	<b>14.8 (10.9)</b>	<b>28.9 (8.0)</b>	<b>90.5 (10.1)</b>	<b>7.0 (7.3)</b>	<b>29.5 (8.5)</b>	<b>91.4 (9.2)</b>	<b>6.4 (7.2)</b>	<b>26.9 (8.7)</b>	<b>90.5 (9.9)</b>	<b>6.8 (8.3)</b>			

Third, we evaluated the relevance of our GAN-based regularization by training our framework with and without ( $\lambda_{GAN} = 0$ ) it on the cross-modal image synthesis. Quantitative and qualitative assessments were performed.

#### 4.4.3 Brain tumor segmentation in iUS

To further compare our proposed framework with competing methods, we conducted another set of experiments on the challenging downstream task of brain tumor segmentation in ultrasound images. While automatic segmentation of brain tumors in iUS could provide intraoperative guidance during surgery, annotating ultrasound images requires time and rare clinical expertise as tumor boundaries are often unclear and ambiguous. For this reason, **only a very small annotated iUS dataset RESECT-SEG (N=22 annotated 3D iUS) is currently available.** In contrast, **the performance of deep learning models for brain tumor segmentation in MR images has reached comparable performance to human annotators [42]** thanks to the availability of large datasets, such as UPenn-

GBM (N=611). In this task, we propose to leverage our cross-modal synthesis framework to learn to perform brain tumor segmentation in iUS by synthesizing iUS from MR data and exploiting available MR annotations, i.e., performing cross-modality domain adaptation for image segmentation [49].

First, **we generate virtual 3D ultrasound sweeps in pre-operative MR data from UPenn-GBM using [50].** Second, we exploit MM-GAN, ResViT, M2DN, and Ours from the synthesis task to synthesize iUS data from 2D MR slices in the virtual sweeps' field of view. This created a paired dataset of synthetic ultrasound images and brain tumor annotations for each framework. Third, we trained a 5-fold nnU-Net ensemble [51], a well-established framework for 3D image segmentation, to perform automated segmentation in the ultrasound images using each synthetic paired dataset. This led to the creation of one ensemble for each cross-modal synthesis technique. To further benchmark performance against a fully supervised baseline, we also trained a 5-fold nnU-Net ensemble on the only publicly available dataset (RESECT-

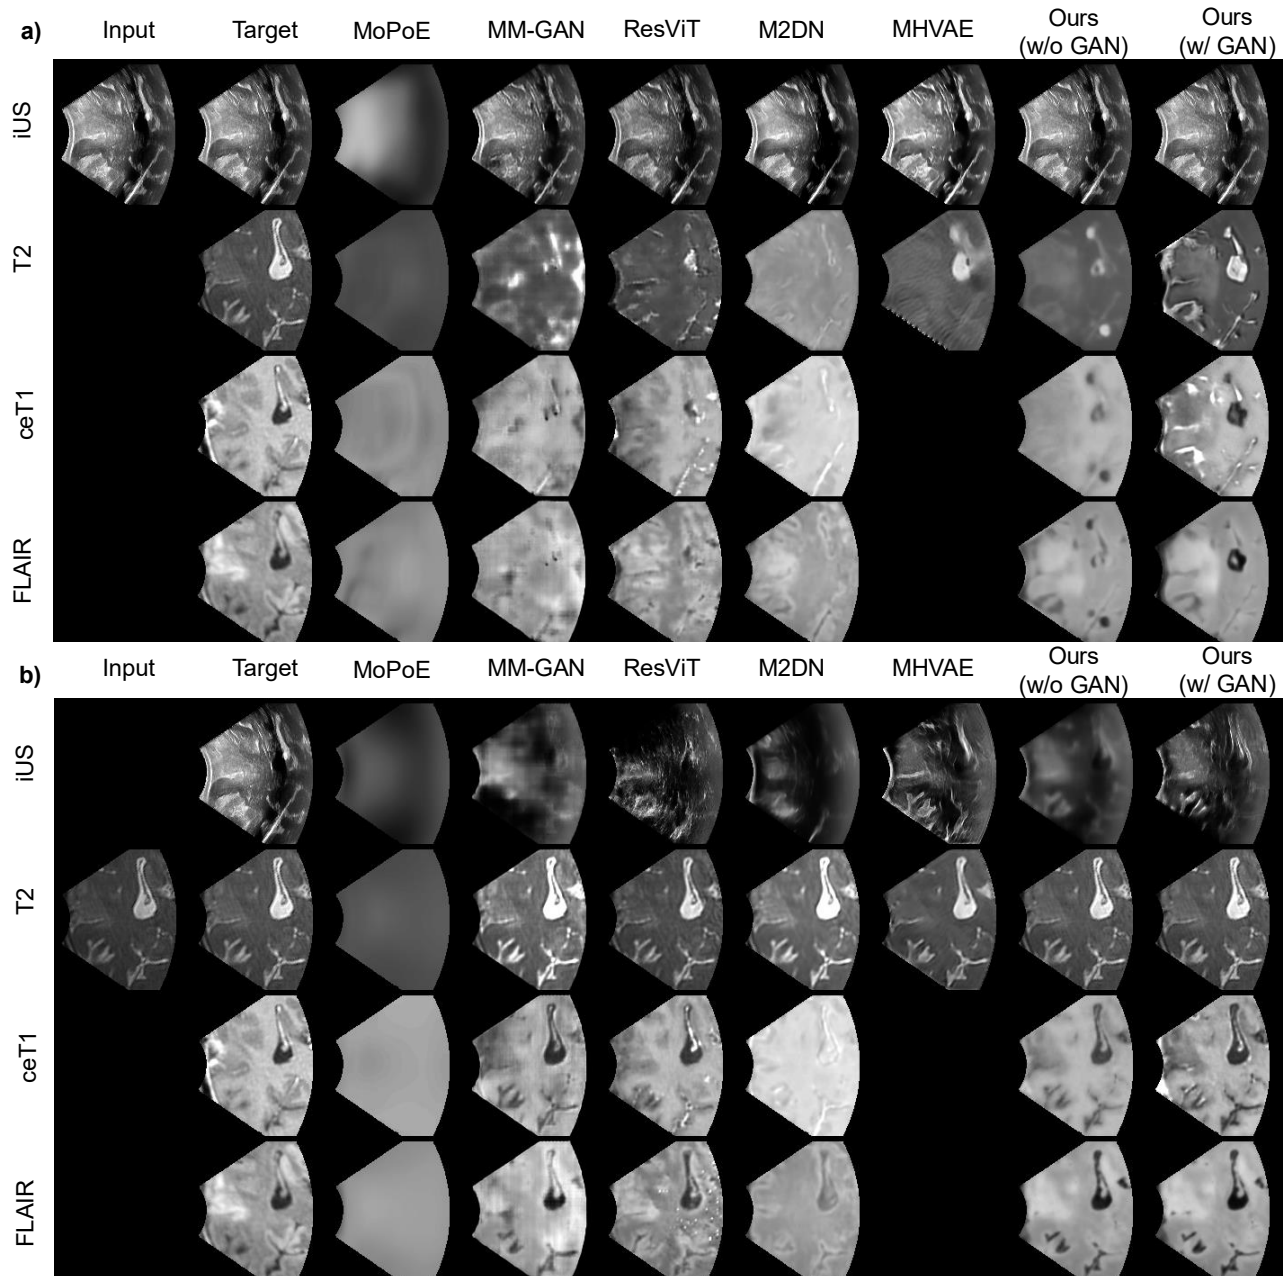


Fig. 3. Qualitative comparison of our method with all competing methods for synthesizing all modalities (iUS,  $T_2$ ,  $ceT_1$ , FLAIR) from (a) iUS; (b)  $T_2$ . Our approach generates sharper images with better contrast differentiation between tissues and modality-specific patterns (e.g. speckles for iUS).

SEG) containing real iUS with tumor annotations. Finally, we assessed the performance of the trained ensembles using the RESECT-SEG dataset with  $N = 22$  annotated ultrasound images and on a subset of  $N=6$  ReMIND iUS images not used during training [52]. Dice Score Coefficient (DSC) and Average Symmetric Surface Distance (ASSD) were computed to compare manual annotations with predictions.

#### 4.4.4 MR-iUS image registration

Finally, we conducted experiments on the problem of multimodal image registration between pre-operative MR and 3D iUS images. Registering MRI and iUS is challenging because they provide different characterizations of tissues associated with different physical properties. Consequently, MRI and iUS provide complementary information in very

different contrasts. To facilitate the registration process, we investigate using our cross-modal synthesis framework. Specifically, we propose registering each pre-operative MRI with its corresponding synthesized sequence from iUS. As we have access to paired data, we simulated rigid deformations within four ranges with median displacements of 0-4 mm, 4-8 mm, 8-12 mm, and 12-16 mm. Then, registration was performed using either the original iUS volume or the simulated MR sequence in the field of view of the iUS using MM-GAN, ResViT, M2DN, and Ours. Gradient-descent registration was used to optimize the 6 degrees of freedom using a multiscale normalized local cross-correlation metric [53]. We reported the average Target Registration Error (TRE) in mm for each MR sequence by comparing true and predicted displacement at each voxel.

## 5 EXPERIMENTAL RESULTS

This section presents the results of the harmonized cross-modal synthesis task, ablation study and downstream tasks.

### 5.1 Harmonized cross-modal image synthesis results

We first assessed the effectiveness of MMHVAE in performing harmonized cross-modal image synthesis of multi-sequence MRI and ultrasound. MMHVAE was compared against the state-of-the-art unified synthesis framework using convolutional (MoPoE, MM-GAN), transformer (ResViT), and diffusion models (M2DN). Performance on unified models was evaluated at test time for each possible combination of input images in the ReMIND dataset. Qualitative results are shown in Figure 3, and detailed quantitative results are presented in Table 1.

**Harmonization task:** When the target MR modality is present in the input set (indicated by a bullet point), the task is to perform image harmonization. Our method consistently outperforms the competing models across all target MR sequence modalities ( $T_2$ ,  $ceT_1$ , and FLAIR). For instance, when  $T_2$  harmonization is performed (non-harmonized  $T_2$  present in input), **our method achieves an SSIM greater than 98.2% and a lower LPIPS than 1.3% for any combination**, indicating near-perfect structural similarity and perceptual quality. In contrast, MoPoE, MM-GAN, ResViT and M2DN achieve lower SSIM scores (82.9%, 93.0%, 93.5%, and 95.0% respectively) and higher LPIPS values (24.3%, 5.4%, 5.6%, and 4.6% respectively). Similar conclusions can be drawn for  $ceT_1$  and FLAIR modalities, where our approach achieves the highest PSNR, SSIM, and the lowest LPIPS. Finally, our approach also significantly outperforms our previously proposed approach, MHVAE, on  $T_2$  harmonization (98.5% vs 91.0% SSIM; 5.7 vs 1.1 LPIPS). This set of experiences demonstrates the superiority of our method in harmonizing images with minimal loss of information.

**Cross-modal synthesis task:** When the target modality is missing in the input set (indicated by an empty circle), the task synthesizes the harmonized missing modality using the available input images. Figure 3 shows qualitative results for various input combinations. Our method consistently outperforms state-of-the-art approaches across all modalities. First, as illustrated in Figure 3, MoPoE produces noticeably blurry images, highlighting its limitations in generating high-quality outputs. In contrast, our approach produces more realistic and detailed images than other unified synthesis methods. Specifically, our method generates synthetic ultrasound images with more realistic textures, such as the presence of speckles, while also better preserving anatomical structures. These qualitative improvements align with the quantitative metrics shown in Table 1, where our method achieves lower perceptual LPIPS and higher structural SSIM scores.

Although MM-GAN, ResViT, and M2DN were primarily designed for synthesizing multi-parametric MR scans, our approach proves more effective at performing cross-modal synthesis between different MR sequences. For instance, our method reconstructs FLAIR scans more accurately from  $ceT_1$ ,  $T_2$ , or a combination of both, compared to these two methods. This shows that while our model learns a

shared latent representation, it outperforms state-of-the-art methods that rely on advanced network architectures.

**Computational resources:** Finally, **our approach is significantly more computationally efficient** than the best-performing baseline, ResViT, with a much lower time complexity (13G vs. 274G MACs) and a smaller model size (14M vs. 293M parameters). Compared to M2DN, which requires 18 inference steps, our method is both faster (55 ms vs. 290 ms) and more efficient (13G vs. 24G MACs). Moreover, our hierarchical fusion strategy introduces only a marginal overhead compared to MoPoE, with a 19% increase in time complexity and a 4% increase in model size.

### 5.2 Ablation study results

#### 5.2.1 Number of levels in the hierarchy

The results of the ablation study, presented in Table 2, clearly demonstrate the positive impact of increasing the number of latent variables  $L$  in the hierarchical representation on the performance of our cross-modal image synthesis framework. **As  $L$  increases from 1 to 7, there is a consistent improvement in the quality of the synthesized images** across all modalities (iUS,  $T_2$ ,  $ceT_1$ , and FLAIR). For instance, with  $L = 2$ , there is already an improvement over the non-hierarchical approach ( $L = 1$ ), with an increase in PSNR values across modalities, particularly for  $T_2$  (22.6 dB for  $L = 2$  vs. 21.6 dB for  $L = 1$ ) and FLAIR (21.1 dB for  $L = 2$  vs. 19.2 dB for  $L = 1$ ). As the number of latent variables increases, these improvements become evident. For example, with  $L = 5$ , the PSNR for  $T_2$  and FLAIR increases to 27.8 dB and 26.8 dB.

The model with  $L = 7$  latent variables achieves the highest performance across all metrics, with a PSNR of 27.0 dB for iUS, 28.6 dB for  $T_2$ , 29.8 dB for  $ceT_1$ , and 27.4 dB for FLAIR. Similarly, the SSIM scores are also highest and the LPIPS values are the lowest, indicating better structural similarity and perceptual quality. In short, as the latent structure becomes more complex with additional variables, the model gains in expressiveness, leading to more accurate and realistic image synthesis.

#### 5.2.2 Fusing operation

The ablation study results in Table 2 demonstrate the effectiveness of our proposed fusion operation compared to two commonly used techniques for handling missing data: **Concat. w/ zeros** and **Average**.

The **Concat. w/ zeros** method, which concatenates available images with zero imputation for missing modalities, obtained lower performance, especially in terms of perceptual quality. For instance, it achieves a PSNR of 26.1 dB, an LPIPS of 15.6% for iUS, and a PSNR of 24.0 dB with an LPIPS of 8.7% for FLAIR.

The **Average** method, which averages features extracted independently from each available image at each level of the hierarchy, performs slightly better, especially in terms of SSIM and PSNR across modalities. For example, it achieves a PSNR of 27.0 dB and an SSIM of 89.0% for  $T_2$ , indicating that extracting features from each image independently helps retain more information across missing modalities.

Our proposed fusion operation outperforms both methods across almost all metrics and modalities. For

TABLE 2

**Ablation studies:** a) effect of number of latent variables, b) fusing operation, and c) GAN loss. Scores are averaged across all possible combinations of inputs. Values in bold indicate the best performance based on a paired Wilcoxon signed-rank test with Bonferroni correction ( $p < 0.01$ ). Experiments performed on one fold of the ReMIND dataset.

Setting	iUS			$T_2$			$ceT_1$			FLAIR		
	PSNR(dB) $\uparrow$	SSIM(%) $\uparrow$	LPIPS(%) $\downarrow$	PSNR(dB) $\uparrow$	SSIM(%) $\uparrow$	LPIPS(%) $\downarrow$	PSNR(dB) $\uparrow$	SSIM(%) $\uparrow$	LPIPS(%) $\downarrow$	PSNR(dB) $\uparrow$	SSIM(%) $\uparrow$	LPIPS(%) $\downarrow$
<b>a) Number of latent variables <math>L</math></b>												
$L = 1$	21.5 (4.5)	75.6 (12.3)	30.7 (13.8)	21.6 (4.9)	83.4 (8.3)	24.1 (10.5)	24.0 (4.4)	86.0 (6.9)	19.2 (9.1)	19.2 (4.6)	80.4 (9.9)	22.9 (10.7)
$L = 2$	21.9 (4.3)	75.5 (11.9)	21.4 (10.2)	22.6 (4.5)	81.4 (9.4)	15.2 (6.7)	24.6 (4.1)	84.0 (7.9)	11.8 (5.8)	21.1 (4.5)	80.9 (9.7)	16.4 (7.8)
$L = 3$	22.8 (4.5)	77.3 (11.2)	19.3 (10.3)	24.6 (4.8)	85.0 (8.3)	11.3 (5.8)	25.7 (4.2)	87.0 (7.4)	9.7 (5.0)	23.7 (4.7)	85.1 (8.1)	10.9 (6.0)
$L = 5$	23.5 (6.1)	78.4 (13.4)	15.0 (8.2)	27.8 (6.9)	89.8 (9.3)	8.0 (6.4)	28.5 (6.4)	90.6 (8.5)	6.5 (5.9)	26.8 (6.6)	89.4 (9.6)	<b>6.5 (5.7)</b>
$L = 7$	<b>27.0 (11.5)</b>	<b>80.6 (15.0)</b>	<b>14.7 (10.9)</b>	<b>28.6 (7.8)</b>	<b>90.7 (9.7)</b>	<b>7.1 (7.1)</b>	<b>29.8 (7.6)</b>	<b>91.7 (8.9)</b>	<b>6.2 (6.3)</b>	<b>27.4 (7.0)</b>	<b>91.0 (9.3)</b>	6.7 (7.2)
<b>b) Fusing operation</b>												
Concat. w/ zeros	26.1 (9.6)	<b>80.7 (14.8)</b>	15.6 (11.2)	25.8 (5.9)	88.6 (10.4)	9.0 (7.1)	26.3 (5.2)	90.1 (8.8)	8.0 (6.0)	24.0 (4.8)	88.2 (9.6)	8.7 (7.0)
Average	26.3 (10.1)	80.2 (15.2)	<b>14.4 (10.2)</b>	27.0 (7.2)	89.0 (10.6)	8.7 (7.8)	27.2 (6.0)	90.0 (9.1)	7.6 (6.4)	25.7 (6.2)	89.7 (9.5)	7.5 (7.3)
Ours	<b>27.0 (11.5)</b>	<b>80.6 (15.0)</b>	<b>14.7 (10.9)</b>	<b>28.6 (7.8)</b>	<b>90.7 (9.7)</b>	<b>7.1 (7.1)</b>	<b>29.8 (7.6)</b>	<b>91.7 (8.9)</b>	<b>6.2 (6.3)</b>	<b>27.4 (7.0)</b>	<b>91.0 (9.3)</b>	<b>6.7 (7.2)</b>
<b>c) GAN loss</b>												
Without GAN	<b>28.1 (12.0)</b>	<b>81.4 (14.5)</b>	20.6 (15.3)	<b>30.1 (8.0)</b>	<b>91.6 (9.0)</b>	8.3 (9.5)	<b>30.6 (8.3)</b>	<b>92.6 (7.9)</b>	7.4 (8.8)	<b>27.2 (8.8)</b>	<b>90.5 (9.8)</b>	<b>6.6 (8.3)</b>
With GAN	27.0 (11.5)	79.7 (15.6)	<b>14.6 (10.9)</b>	28.8 (8.0)	89.8 (10.8)	<b>7.1 (7.4)</b>	29.4 (8.4)	90.6 (9.9)	<b>6.3 (7.0)</b>	26.8 (8.6)	90.1 (10.2)	<b>6.6 (8.1)</b>

TABLE 3

**Comparison against the state-of-the-art unified models on the image registration downstream task** between iUS and rigidly deformed (a)  $T_2$ , (b)  $ceT_1$  and (c) FLAIR scans. The registration algorithm either uses the acquired iUS as a reference image or the translated iUS in the modality of the moving image. Best performance in bold is based on a paired Wilcoxon signed-rank test ( $p < 0.05$ ). Mean and standard deviation of Target Registration Error are given in mm for four ranges of displacements.

	<b>(a) iUS to <math>T_2</math> registration</b>					<b>(b) iUS to <math>ceT_1</math> registration</b>					<b>(c) iUS to FLAIR registration</b>				
	0-4 mm	4-8 mm	8-12 mm	12-16 mm		0-4 mm	4-8 mm	8-12 mm	12-16 mm		0-4 mm	4-8 mm	8-12 mm	12-16 mm	
Acquired iUS	1.0 (0.9)	1.5 (1.7)	2.0 (2.7)	7.3 (7.3)		12.6 (3.2)	12.9 (3.7)	14.6 (3.8)	14.9 (4.0)		8.9 (2.6)	9.2 (3.0)	9.4 (3.2)	12.6 (4.4)	
Tran. $T_2$	MM-GAN [14]	1.2 (0.8)	2.0 (2.5)	4.4 (5.3)	9.6 (7.3)	MM-GAN [14]	6.3 (3.7)	7.7 (5.2)	7.6 (5.1)	11.0 (5.2)	MM-GAN [14]	7.5 (4.1)	8.5 (4.0)	10.4 (4.1)	12.2 (4.3)
	ResViT [15]	<b>0.8 (0.3)</b>	<b>1.4 (2.4)</b>	<b>2.5 (4.0)</b>	<b>7.1 (7.2)</b>	ResViT [15]	<b>2.2 (2.4)</b>	<b>4.0 (4.3)</b>	<b>4.1 (4.3)</b>	<b>7.9 (5.7)</b>	ResViT [15]	<b>1.9 (2.7)</b>	<b>3.4 (4.6)</b>	<b>5.1 (5.0)</b>	10.8 (6.2)
	M2DN [36]	1.1 (0.7)	3.5 (4.4)	5.8 (5.3)	<b>9.2 (6.1)</b>	M2DN [36]	5.4 (2.9)	6.6 (3.8)	9.8 (4.4)	12.2 (4.0)	M2DN [36]	4.6 (3.5)	7.1 (4.7)	9.4 (4.7)	13.2 (5.2)
	Ours	<b>0.7 (0.2)</b>	<b>1.3 (2.3)</b>	<b>2.0 (3.4)</b>	<b>7.3 (6.9)</b>	Ours	<b>2.3 (1.7)</b>	<b>2.8 (2.4)</b>	<b>3.2 (3.3)</b>	<b>6.5 (5.0)</b>	Ours	<b>1.8 (1.2)</b>	<b>2.3 (2.1)</b>	<b>4.4 (4.7)</b>	<b>8.0 (6.5)</b>

instance, it achieves a PSNR of 27.0 dB for iUS, 28.6 dB for  $T_2$ , 29.8 dB for  $ceT_1$ , and 27.4 dB for FLAIR. The SSIM scores are also the highest (e.g., 91.7% for  $ceT_1$  and 91.0% for FLAIR), while the LPIPS scores are the lowest. These results show that our probabilistic fusion approach, which integrates information more effectively across available modalities, leads to more accurate and visually realistic image synthesis, particularly compared to more straightforward concatenation or averaging methods.

### 5.2.3 GAN regularization

Another key component of our framework is the use of the GAN loss. We present results with ( $\lambda_{GAN} = 0.025$ ) and without ( $\lambda_{GAN} = 0$ ) the GAN loss in Figure 3 and Table 2. As shown in Figure 3, using the GAN loss improves the visual quality of the generated images. Without the GAN loss, our model tends to produce smoother and less realistic images, particularly lacking the speckles characteristic of ultrasound images. These qualitative improvements align with the perceptual results in Table 2, where lower LPIPS scores are achieved with the GAN loss. However, we also found that the GAN loss leads to lower SSIM and PSNR scores. This illustrates the typical challenge in assessing image synthesis methods, where metrics may not agree.

This challenge motivated us to conduct additional experiments to further evaluate the quality of the synthesized images in two downstream tasks, where errors can be quantified with actionable and interpretable metrics: image segmentation and registration.

TABLE 4

**Comparison against the state-of-the-art unified models on brain tumor segmentation downstream tasks in iUS.** Best performance in bold is based on a paired Wilcoxon signed-rank test ( $p < 0.05$ ).

	RESECT-SEG [43]		ReMIND [1]	
	Dice Score (%) $\uparrow$	ASSD (mm) $\downarrow$	Dice Score (%) $\uparrow$	ASSD (mm) $\downarrow$
MM-GAN [14]	53.3 [37.1 - 69.8]	4.1 [3.4 - 5.3]	57.9 [52.6 - 71.5]	5.0 [3.3 - 6.0]
ResViT [15]	<b>67.3 [48.3 - 80.5]</b>	<b>2.3 [1.7 - 3.7]</b>	<b>74.7 [69.0 - 79.6]</b>	<b>2.6 [2.3 - 3.3]</b>
M2DN [36]	65.2 [49.2 - 76.6]	3.0 [2.4 - 4.6]	62.2 [49.2 - 63.3]	4.6 [4.2 - 5.7]
Ours	<b>73.6 [54.4 - 81.3]</b>	<b>2.3 [1.5 - 4.0]</b>	<b>77.6 [67.6 - 84.4]</b>	<b>2.4 [1.7 - 3.6]</b>
Fully [43]	-	-	<b>73.4 [67.5 - 75.7]</b>	<b>2.4 [2.2 - 2.4]</b>
Expert	-	-	84.2 [83.3 - 84.8]	1.5 [1.0 - 1.6]

### 5.3 Learning brain tumor Segmentation in iUS from synthetic iUS

The results of the downstream task of brain tumor segmentation in ultrasound images are presented in Table 4. All nnU-Net models were trained using synthetic ultrasound images generated by each cross-modal synthesis method from the same MR dataset. This MR dataset contains all combinations of possible input MR data ( $T_1$ ,  $T_2$ , FLAIR,  $T_1+T_2$ , etc). Our proposed framework demonstrated superior performance at synthesizing ultrasound data compared to competing approaches, achieving the highest Dice scores on both the RESECT-SEG (73.6%) and ReMIND (77.6%) datasets and significantly outperforming MM-GAN and M2DN. These results show that the ultrasound images synthesized by our method from MR data are accurate, particularly in regions in and near the tumor. Notably, our

approach reached comparable performance to an ensemble of nnU-Net models trained using a small annotated dataset of real iUS, achieving a higher median Dice score (77.6% vs 73.4%), while maintaining the same median ASSD (2.4 mm). This highlights not only the quality of the generated images but also their utility for training segmentation models in a new modality, such as ultrasound, without requiring additional annotations. Finally, our approach achieves performance close to that of expert neurosurgeon, while enabling rapid 3D iUS delineation in under 10 seconds compared to experienced neurosurgeons (32-93 minutes).

#### 5.4 Improving iUS-MR image registration with synthetic MRI

The results of the multimodal image registration task between MRI and 3D intraoperative ultrasound (iUS) images are presented in Table 3, where various approaches are compared in terms of Target Registration Error (TRE).

We found that performing image registration using synthetic MR data generated from iUS, rather than acquired 3D iUS data, significantly improves the registration process. Notably, the complexity of registering iUS with MRI varies depending on the MR sequence. The registration performs well with  $T_2$  images even with large initial displacements (up to 12 mm). In contrast, for  $ceT_1$  or FLAIR sequences, the registration often converges to incorrect solutions ( $>10$  mm), even for small initial displacements (0-4 mm). Conversely, registering synthetic MR image from our method leads to significant improvements consistently for all MR sequences, even with large initial displacements (up to 12 mm). For instance, the average registration error is reduced from 14.6 to 3.2 mm for  $ceT_1$  in the 8-12 mm displacement range. While our approach does not always fully converge to correct solutions for substantial displacements (12-16 mm), it still reduces the displacement by a factor of two on average. Overall, these findings demonstrate that synthesizing MR images from iUS significantly improves registration accuracy, particularly for  $ceT_1$  and FLAIR sequences and larger deformations, where contrast differences often cause conventional registration methods to converge to suboptimal solutions.

We also compared our framework with other synthesis frameworks (MM-GAN, ResViT, and M2DN). Our method consistently outperformed the competing approaches in terms of TRE, particularly for larger displacement ranges. For instance, in the  $T_2$  registration, our approach obtained the lowest TRE across all ranges, achieving 2.0 mm for the 8-12 mm range, compared to 4.4 mm for MM-GAN, 2.5 mm for ResViT, and 5.8 mm for M2DN. Similarly, for  $ceT_1$ , our method achieved a TRE of 2.8 mm for the 4-8 mm range, significantly outperforming MM-GAN (7.7 mm), ResViT (4.0 mm), and M2DN (6.6 mm). Our method’s advantage was even more evident for the FLAIR modality, with a TRE of 1.8 mm for the 0-4 mm range, compared to 7.5 mm, 1.9 mm, and 4.6 mm for MM-GAN, ResViT, and M2DN, respectively. Overall, these results show our translated MR images better preserve anatomical structures in a similar contrast to the target MR images.

## 6 DISCUSSION AND CONCLUSION

In this section, we discuss some additional components of our framework and conclude this work.

### 6.1 Latent Representation Analysis

This section analyzes the learned latent representation at the pixel level. A key limitation of MVAEs [23] is the misalignment of latent representations, where the model tends to learn inconsistent representations for different input data modalities. To address this, we modeled the variational posterior as a mixture of product-of-experts. As shown in Appendix 1.2, this mixture encourages the approximate posterior distribution from incomplete data to align with the true posterior from complete data, thereby estimating the missing information and ensuring consistency across the different input modalities.

To experimentally assess the alignment of the latent representations, we performed a Principal Component Analysis (PCA) on the latent variable  $z_1$  from various input combinations (iUS,  $T_2$ ,  $ceT_1$ , or FLAIR) and from all modalities combined. The PCA was conducted by randomly selecting  $N = 10,000$  pixel representations from all slices of a patient and all combinations of input data. The results, presented in Figure 4, show a good alignment between the latent representations across different modalities despite the inherent differences in contrast and information between them. Interestingly, the first principal component of the  $z_1$  obtained using all images presents a clear contrast between ventricles and white matter, as in a  $T_2$  image, and clear delineation of the tumor, as in a  $ceT_1$  image. While this information is not present in the iUS, its latent representation obtained from iUS tries to estimate it. These results highlight the effectiveness of our mixture of product-of-experts approach, as it successfully produces well-aligned latent representations for different inputs.

### 6.2 Variability in the samples

It is common to lower the temperature of the conditional distributions when sampling from HVAEs on challenging datasets [25]. This is done by scaling down the standard deviation of the Normal distributions at each level in the approximate posterior. It often improves the samples’ quality but also reduces their diversity. Figure 5 shows how the cross-modal synthesis quality and diversity vary with temperature. We found that sampling with a low temperature tends to produce smooth images that accurately reflect the overall structure of the images. However, these images are less realistic due to their lack of finer, modality-specific details. In contrast, higher temperatures introduce modality-specific features, such as speckles in synthetic ultrasound from  $T_2$  or contrast agents in contrast-enhanced  $T_1$  images from iUS, which correspond to information not explicitly present in the conditioning image. This indicates that modality-specific details are captured in the variability of the Normal distribution. By adjusting the temperature, we can modulate how much of this modality-specific information is incorporated into the synthesized images, with higher temperatures allowing for the generation of more realistic but less reliable images. We




Fig. 4. Principal Component Analysis on the first three components (PC1: 39%, PC2: 16%, PC3: 13%) of the latent variable at the highest level  $z_1$  estimated using (a)  $iUS+T_2+ceT_1+FLAIR$ ; (b)  $iUS$ ; (c)  $T_2$  (d)  $ceT_1$  (e)  $FLAIR$  as input. Similar representations are obtained for all combinations, in particular in the tumor region (red) and around the ventricle (blue).




Fig. 5. Impact of the temperature  $T$  on the quality of the reconstructed images for a)  $iUS$  to  $T_2$ ; b)  $iUS$  to  $ceT_1$  synthesis.

found that a temperature of 0.5 was a good compromise for the realism and variability of samples.

### 6.3 Limitations

While MMHVAE demonstrates promising results for medical multimodal synthesis, several limitations should be acknowledged. First, the framework has been primarily evaluated on imaging datasets, and its applicability to non-imaging modalities (e.g., genomics or clinical data) remains to be established. Second, the method relies on the availability of sufficiently paired multimodal data during training; in extreme cases where a modality is consistently missing or available in isolation, performance is expected to degrade. Finally, the interpretability of the learned latent representations remains limited, and future work is needed to link the latent factors to clinically meaningful variables.

### 6.4 Conclusion

We proposed a framework for synthesizing missing imaging data across modality, with the flexibility to handle all combination of data missingness at training and inference time. The proposed model relies on a mixture of product-of-experts to encode observed information and estimate missing ones. Our experiment results verified the effectiveness of our model empirically.

### ACKNOWLEDGMENTS

This work was supported by the National Institutes of Health (R01EB032387, R01EB027134, P41EB015902, P41EB028741 and K25EB035166). R.D. received a Marie Skłodowska-Curie grant No 101154248 (project: SafeREG). This work was performed using HPC resources from GENCI-IDRIS (Grant 2024-SafeREG, 2025-SafeREG). The research leading to these results has received funding from the program "Investissements d'avenir" ANR-10-IAIHU-06.

## REFERENCES

- [1] P. Juvekar, R. Dorent, F. Kögl, E. Torio, C. Barr, L. Rigolo, C. Galvin, N. Jowkar, A. Kazi, N. Haouchine *et al.*, "Remind: The brain resection multimodal imaging database," *Scientific Data*, vol. 11, no. 1, p. 494, 2024.
- [2] H. Shen, X. Li, Q. Cheng, C. Zeng, G. Yang, H. Li, and L. Zhang, "Missing Information Reconstruction of Remote Sensing Data: A Technical Review," *IEEE Geoscience and Remote Sensing Magazine*, vol. 3, no. 3, pp. 61–85, 2015.
- [3] Z. Wang, Y. Wu, and Q. Niu, "Multi-Sensor Fusion in Automated Driving: A Survey," *IEEE Access*, vol. 8, pp. 2847–2868, 2020.
- [4] R. J. Little and D. B. Rubin, *Statistical analysis with missing data*. John Wiley & Sons, 2019, vol. 793.
- [5] D. B. Rubin, "Multiple Imputation After 18+ Years," *Journal of the American Statistical Association*, vol. 91, no. 434, pp. 473–489, 1996.
- [6] O. Troyanskaya, M. Cantor, G. Sherlock, P. Brown, T. Hastie, R. Tibshirani, D. Botstein, and R. B. Altman, "Missing value estimation methods for DNA microarrays," *Bioinformatics*, vol. 17, no. 6, pp. 520–525, 06 2001.
- [7] Y. Dong, Q.-A. Fu, X. Yang, T. Pang, H. Su, Z. Xiao, and J. Zhu, "Benchmarking adversarial robustness on image classification," in *CVPR*, June 2020.
- [8] T. Wang, Y. Lei, Y. Fu, J. F. Wynne, W. J. Curran, T. Liu, and X. Yang, "A review on medical imaging synthesis using deep learning and its clinical applications," *Journal of applied clinical medical physics*, vol. 22, no. 1, pp. 11–36, 2021.
- [9] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-To-Image Translation With Conditional Adversarial Networks," in *CVPR*, 2017.
- [10] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic Image Synthesis with Spatially-Adaptive Normalization," in *CVPR*, 2019.
- [11] M. Donnez, F.-X. Carton, F. Le Lann, E. De Schlichting, and M. Chabanas, "Realistic synthesis of brain tumor resection ultrasound images with a generative adversarial network," in *Medical Imaging 2021: Image-Guided Procedures, Robotic Interventions, and Modeling*, vol. 11598. SPIE, 2021, pp. 637–642.
- [12] J. Jiao, A. I. L. Namburete, A. T. Papageorghiou, and J. A. Noble, "Self-Supervised Ultrasound to MRI Fetal Brain Image Synthesis," *IEEE Transactions on Medical Imaging*, vol. 39, no. 12, 2020.
- [13] A. Chartsias, T. Joyce, G. Papanastasiou, S. Semple, M. Williams, D. E. Newby, R. Dharmakumar, and S. A. Tsaftaris, "Disentangled representation learning in cardiac image analysis," *Medical image analysis*, vol. 58, p. 101535, 2019.
- [14] A. Sharma and G. Hamarneh, "Missing MRI Pulse Sequence Synthesis Using Multi-Modal Generative Adversarial Network," *IEEE Transactions on Medical Imaging*, vol. 39, no. 4, pp. 1170–1183, 2020.
- [15] O. Dalmaz, M. Yurt, and T. Çukur, "Resvit: Residual vision transformers for multimodal medical image synthesis," *IEEE Transactions on Medical Imaging*, vol. 41, no. 10, pp. 2598–2614, 2022.
- [16] H. Li, J. C. Paetzold, A. Sekuboyina, F. Kofler, J. Zhang, J. S. Kirschke, B. Wiestler, and B. Menze, "DiamondGAN: Unified Multi-Modal Generative Adversarial Networks for MRI Sequences Synthesis," in *MICCAI 2019*, 2019.
- [17] D. Lee, J. Kim, W.-J. Moon, and J. C. Ye, "CollaGAN: Collaborative GAN for missing image data imputation," in *CVPR*, 2019, pp. 2487–2496.
- [18] M. Havaei, N. Guizard, N. Chapados, and Y. Bengio, "Hemis: Hetero-modal image segmentation," in *MICCAI 2016*. Cham: Springer, 2016, pp. 469–477.
- [19] T. Varsavsky and *et al.*, "PIMMS: permutation invariant multi-modal segmentation," in *DLMIA 2018, MICCAI 2018*, 2018, pp. 201–209.
- [20] R. Dorent, T. Booth, W. Li, C. H. Sudre, S. Kafiabadi, J. Cardoso, S. Ourselin, and T. Vercauteren, "Learning joint segmentation of tissues and brain lesions from task-specific hetero-modal domain-shifted datasets," *Medical image analysis*, vol. 67, 2021.
- [21] A. Chartsias, T. Joyce, M. V. Giuffrida, and S. A. Tsaftaris, "Multimodal mr synthesis via modality-invariant latent representation," *IEEE transactions on medical imaging*, vol. 37, no. 3, pp. 803–814, 2017.
- [22] T. Zhou, H. Fu, G. Chen, J. Shen, and L. Shao, "Hi-net: hybrid-fusion network for multi-modal mr image synthesis," *IEEE transactions on medical imaging*, 2020.
- [23] M. Wu and N. Goodman, "Multimodal Generative Models for Scalable Weakly-Supervised Learning," *NeurIPS*, vol. 31, 2018.
- [24] R. Dorent, S. Joutard, M. Modat, S. Ourselin, and T. Vercauteren, "Hetero-modal variational encoder-decoder for joint modality completion and segmentation," in *MICCAI 2019*, 2019.
- [25] A. Vahdat and J. Kautz, "NVAE: A Deep Hierarchical Variational Autoencoder," *NeurIPS*, vol. 33, 2020.
- [26] R. Ranganath, D. Tran, and D. Blei, "Hierarchical Variational Models," in *ICML*, 2016.
- [27] L. Maaløe, M. Fraccaro, V. Liévin, and O. Winther, "BIVA: A Very Deep Hierarchy of Latent Variables for Generative Modeling," *NeurIPS*, 2019.
- [28] C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby, and O. Winther, "Ladder Variational Autoencoders," *NeurIPS*, 2016.
- [29] R. Dorent, N. Haouchine, F. Kögl, S. Joutard, P. Juvekar, E. Torio, A. J. Golby, S. Ourselin, S. Frisken, T. Vercauteren, T. Kapur, and W. M. Wells, "Unified Brain MR-Ultrasound Synthesis Using Multi-modal Hierarchical Representations," in *MICCAI 2023*, 2023, pp. 448–458.
- [30] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *ICLR*, 2014.
- [31] Y. Shi, B. Paige, P. Torr *et al.*, "Variational mixture-of-experts autoencoders for multi-modal deep generative models," *NeurIPS*, vol. 32, 2019.
- [32] J. Kittler, M. Hatef, R. P. Duin, and J. Matas, "On combining classifiers," *IEEE transactions on pattern analysis and machine intelligence*, vol. 20, no. 3, pp. 226–239, 1998.
- [33] T. M. Sutter, I. Daunhawer, and J. E. Vogt, "Generalized multimodal ELBO," in *ICLR*, 2021.
- [34] L. Jiang, Y. Mao, X. Wang, X. Chen, and C. Li, "Coladiff: Conditional latent diffusion model for multi-modal mri synthesis," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2023, pp. 398–408.
- [35] X. Meng, Y. Gu, Y. Pan, N. Wang, P. Xue, M. Lu, X. He, Y. Zhan, and D. Shen, "A novel unified conditional score-based generative framework for multi-modal medical image completion," *arXiv preprint arXiv:2207.03430*, 2022.
- [36] X. Meng, K. Sun, J. Xu, X. He, and D. Shen, "Multi-modal modality-masked diffusion network for brain mri synthesis with random modality missing," *IEEE Transactions on Medical Imaging*, 2024.
- [37] X. Xiao, Q. V. Hu, and G. Wang, "Fgc2f-udiff: Frequency-guided and coarse-to-fine unified diffusion model for multi-modality missing mri synthesis," *IEEE Transactions on Computational Imaging*, 2024.
- [38] X. Gu, C. Du, T. Pang, C. Li, M. Lin, and Y. Wang, "On memorization in diffusion models," *arXiv preprint arXiv:2310.02664*, 2023.
- [39] M. Modat, G. R. Ridgway, Z. A. Taylor, M. Lehmann, J. Barnes, D. J. Hawkes, N. C. Fox, and S. Ourselin, "Fast free-form deformation using graphics processing units," *Computer Methods and Programs in Biomedicine*, vol. 98, 2010.
- [40] D. Drobny, T. Vercauteren, S. Ourselin, and M. Modat, "Registration of MRI and iUS data to compensate brain shift using a symmetric block-matching based approach," in *CuRIOUS*, 2018.
- [41] S. Bakas, C. Sako, H. Akbari, M. Bilello, A. Sotiras, G. Shukla, J. D. Rudie, N. F. Santamaría, A. F. Kazerooni, S. Pati *et al.*, "The University of Pennsylvania glioblastoma (UPenn-GBM) cohort: advanced MRI, clinical, genomics, & radiomics," *Scientific data*, vol. 9, no. 1, p. 453, 2022.
- [42] S. Bakas and *et al.*, "Identifying the Best Machine Learning Algorithms for Brain Tumor Segmentation, Progression Assessment, and Overall Survival Prediction in the BRATS Challenge," 2019.
- [43] B. Behboodi, F.-X. Carton, M. Chabanas, S. D. Ribaupierre, O. Solheim, B. K. R. Munkvold, H. Rivaz, Y. Xiao, and I. Reinertsen, "RESECT-SEG: Open access annotations of intra-operative brain tumor ultrasound images," 2022. [Online]. Available: <https://arxiv.org/abs/2207.07494>
- [44] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *CVPR*, 2018.
- [45] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *CVPR*, 2018.
- [46] B. Li, K. Xue, B. Liu, and Y.-K. Lai, "Bbdm: Image-to-image translation with brownian bridge diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 1952–1961.

- [47] B. Billot, D. N. Greve, O. Puonti, A. Thielscher, K. Van Leemput, B. Fischl, A. V. Dalca, and J. E. Iglesias, "Synthseg: Segmentation of brain MRI scans of any contrast and resolution without retraining," *Medical Image Analysis*, vol. 86, p. 102789, 2023.
- [48] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric," in *CVPR*, 2018.
- [49] R. Dorent, A. Kujawa, M. Ivory, S. Bakas, N. Rieke *et al.*, "Crossmoda 2021 challenge: Benchmark of cross-modality domain adaptation techniques for vestibular schwannoma and cochlea segmentation," *Medical Image Analysis*, vol. 83, 2023.
- [50] R. Dorent, E. Torio, N. Haouchine, C. Galvin, S. Frisken, A. Golby, T. Kapur, and W. M. Wells, "Patient-specific real-time segmentation in trackerless brain ultrasound," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2024, pp. 477–487.
- [51] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation," *Nature methods*, vol. 18, no. 2, pp. 203–211, 2021.
- [52] R. Dorent, E. Torio, N. Haouchine, P. Juvekar, S. Frisken, A. J. Golby, T. Kapur, and W. Wells, "223 an artificial intelligence framework for brain tumor delineation in intraoperative ultrasound," *Neurosurgery*, vol. 70, no. Supplement\_1, p. 60, 2024.
- [53] T. C. W. Mok and A. C. S. Chung, "Large deformation diffeomorphic image registration with laplacian pyramid networks," in *MICCAI 2020*, 2020, pp. 211–221.




**Sarah Frisken** received a PhD degree from Carnegie Mellon University. She was a Distinguished Research Scientist at Mitsubishi Electric Research Labs, Professor of Computer Science at Tufts University, and CEO of 61 Solutions Inc. She is currently an Associate Professor at Harvard Medical School and Lead Investigator at Brigham and Women's Hospital. Her research interests include shape representation and modeling, medical image processing, and image-guided neurosurgery.



**Reuben Dorent** received a PhD degree from King's College London, in 2022. He then joined the Surgical Planning Laboratory, Brigham and Women's Hospital, Harvard Medical School. He is now a Marie Skłodowska-Curie fellow at Inria Paris-Saclay and the Paris Brain Institute. His research focuses on medical image segmentation, synthesis, cross-modal domain adaptation, and optimization strategies for missing data in medical contexts.



**Tina Kapur** (Member IEEE) received the SM and PhD degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology. She is Associate Professor of Radiology at Harvard Medical School and Brigham and Women's Hospital, where she is Lead Investigator and Executive Director of the Image-Guided Therapy Program. Her research interests include image segmentation, registration, surgical navigation, image-guided therapy, and open-source platforms for medical imaging.



**Nazim Haouchine** received his PhD in Computer Science from the University of Lille and INRIA in 2015, in France. He is now an Assistant Professor at Harvard Medical School and a Research Associate at the Brigham and Women's Hospital. His research interests are in building translational technologies for computer-assisted medical interventions, combining research on computer vision and graphics, deep learning, and physics-based simulation.



**Alexandra J. Golby** received her M.D. from Stanford Medical School. She completed her clinical training in neurosurgery at Stanford and at Brigham and Women's/Boston Children's Hospitals. She is a Neurosurgeon at Brigham and Women's Hospital and a Professor of Neurosurgery and Radiology at Harvard Medical School. Her research focuses on advanced image-guided neurosurgery, functional brain mapping, intra-operative imaging, and minimally invasive approaches.



**William Wells** is Professor of Radiology at Harvard Medical School and Brigham and Women's Hospital, a research scientist at the MIT Computer Science and Artificial Intelligence Laboratory, and a member of the affiliated faculty of the Harvard-MIT division of Health Sciences and Technology. He received a PhD in computer vision from MIT, and has pursued research in medical image computing at the Surgical Planning Laboratory. He is widely known for his ground-breaking work on multi-modality registration by maximization of Mutual Information.




Fig. 6. Intensity distribution of (a)  $T_2$ ; (b)  $ceT_1$ ; (c) FLAIR images using either min-max normalization and the proposed harmonization technique.

## 7 APPENDIX

### 7.1 Proof ELBO $\mathcal{L}_{\text{MMHVAE}}^{\text{ELBO}}$

Let  $\mathbf{x}_r^o$  be an observed set of images. For any distribution  $q_\phi(\mathbf{z}|\mathbf{x}_r^o)$ , the evidence  $\log p_\theta(\mathbf{x}_r^o)$  is lower-bounded by the tractable variational ELBO  $\mathcal{L}^{\text{ELBO}}(\mathbf{x}; \theta, \phi)$ :

$$\mathcal{L}^{\text{ELBO}}(\mathbf{x}_r^o; \theta, \phi) = \sum_{\substack{j=1 \\ \text{s.t. } r_j=1}}^M \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}_r^o)}[\log p_\theta(x_j|\mathbf{z})] - \text{KL}[q_\phi(\mathbf{z}|\mathbf{x}_r^o)||p_\theta(\mathbf{z})]. \quad (20)$$

In this work, we propose to express the variational  $q_\phi(\mathbf{z}|\mathbf{x}_r^o)$  as a mixture (convex combination) of distributions represented by:

$$q_\phi^{\text{MMHVAE}}(\mathbf{z}|\mathbf{x}_r^o) = \sum_{\mathbf{r}' \in S_r} \alpha_{\mathbf{r}'}^{(r)} q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o). \quad (21)$$

Therefore:

$$\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}_r^o)}[\log p_\theta(x_j|\mathbf{z})] = \sum_{\mathbf{r}' \in S_r} \alpha_{\mathbf{r}'}^{(r)} \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)}[\log p_\theta(x_j|\mathbf{z})]. \quad (22)$$

Moreover, as the KL divergence is convex in the pair of probability measures and  $\sum_{\mathbf{r}' \in S_r} \alpha_{\mathbf{r}'}^{(r)} = 1$ :

$$\text{KL}[q_\phi(\mathbf{z}|\mathbf{x}_r^o)||p_\theta(\mathbf{z})] \leq \sum_{\mathbf{r}' \in S_r} \alpha_{\mathbf{r}'}^{(r)} \text{KL}[q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)||p_\theta(\mathbf{z})] \quad (23)$$

Hence, by combining Eq. (20), (22) and (23), the evidence  $\log p(\mathbf{x}_r^o)$  is lower-bounded by:

$$\mathcal{L}_{\text{MMHVAE}}^{\text{ELBO}}(\mathbf{x}_r^o; \theta, \phi) \triangleq \sum_{\mathbf{r}' \in S_r} \alpha_{\mathbf{r}'}^{(r)} \mathcal{L}(\mathbf{x}_{\mathbf{r}'}^o, \mathbf{r}'; \theta, \phi) \quad (24)$$

where:

$$\begin{aligned} \mathcal{L}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi) &= \sum_{\substack{j=1 \\ \text{s.t. } r'_j=1}}^M \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)}[\log p_\theta(x_j|\mathbf{z})] \\ &+ \sum_{\substack{j=1 \\ \text{s.t. } (1-r'_j)r_j=1}}^M \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)}[\log p_\theta(x_j|\mathbf{z})] \\ &- \text{KL}[q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)||p_\theta(\mathbf{z})] \end{aligned} \quad (25)$$

### 7.2 Proof maximal value of variational with respect to $\phi$

Let  $\mathbf{x}_r^o$  be an observed set of images.

$$\begin{aligned} \text{KL}[q_\phi(\mathbf{z}|\mathbf{x}_r^o)||p_\theta(\mathbf{z}|\mathbf{x}_r^o)] &= - \int_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x}_r^o) \log \frac{p_\theta(\mathbf{z}|\mathbf{x}_r^o)}{q_\phi(\mathbf{z}|\mathbf{x}_r^o)} d\mathbf{z} \\ &= - \int_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x}_r^o) \log \frac{p_\theta(\mathbf{z}, \mathbf{x}_r^o)}{q_\phi(\mathbf{z}|\mathbf{x}_r^o) p_\theta(\mathbf{x}_r^o)} d\mathbf{z} \\ &= - \int_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x}_r^o) \log \frac{p_\theta(\mathbf{x}_r^o|\mathbf{z}) p_\theta(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{x}_r^o)} d\mathbf{z} + \log p_\theta(\mathbf{x}_r^o) \\ &= - \mathbb{E}_{\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x}_r^o)} [p_\theta(\mathbf{x}_r^o|\mathbf{z})] + \text{KL}[q_\phi(\mathbf{z}|\mathbf{x}_r^o)||p_\theta(\mathbf{z})] \\ &\quad + \log p_\theta(\mathbf{x}_r^o) \\ &= \log p_\theta(\mathbf{x}_r^o) - \mathcal{L}(\mathbf{x}_r^o, \mathbf{r}'; \theta, \phi) \end{aligned} \quad (26)$$

Consequently, as  $\sum_{\mathbf{r}' \in S_r} \alpha_{\mathbf{r}'}^{(r)} = 1$ , we can find the gap between the evidence  $\log p_\theta(\mathbf{x}_r^o)$  and the ELBO  $\mathcal{L}^{\text{ELBO}}(\mathbf{x}_r^o; \theta, \phi)$ :

$$\begin{aligned} \log p_\theta(\mathbf{x}_r^o) - \mathcal{L}^{\text{ELBO}}(\mathbf{x}_r^o; \theta, \phi) &= \sum_{\mathbf{r}' \in S_r} \alpha_{\mathbf{r}'}^{(r)} \text{KL}[q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)||p_\theta(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o)] \end{aligned} \quad (27)$$

Consequently, as  $\alpha_{\mathbf{r}'}^{(r)} \geq 0$  the maximal value of the ELBO  $\mathcal{L}^{\text{ELBO}}(\mathbf{x}_r^o; \theta, \phi)$  with respect to  $\phi$  is reached when:

$$\forall \mathbf{r}' \in S_r, q_\phi(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o) = p_\theta(\mathbf{z}|\mathbf{x}_{\mathbf{r}'}^o) \quad (28)$$

### 7.3 Proof posterior parameterization

Using Bayes' rule, the true conditional distribution  $p_\theta(\mathbf{z}|\mathbf{x}_r^o)$  can be written as:  $p_\theta(\mathbf{z}|\mathbf{x}_r^o) = \frac{p_\theta(\mathbf{z})}{p(\mathbf{x}_r^o)} p(\mathbf{x}_r^o|\mathbf{z})$ , where  $p(\mathbf{x}_r^o|\mathbf{z}) = \prod_{j=1}^M \text{s.t. } r_j=1 p(x_j|\mathbf{z})$  (Eq. 4). Consequently:

$$p_\theta(\mathbf{z}|\mathbf{x}_r^o) \propto \prod_{\substack{j=1 \\ \text{s.t. } r_j=1}}^M p(x_j|\mathbf{z}) \quad (29)$$

Then, the marginal likelihoods  $p_\theta(x_j|\mathbf{z})$  can be factorized as:

$$\begin{aligned}
p_\theta(x_j|\mathbf{z}) &= \frac{p_\theta(x_j, z_1, \dots, z_L)}{p_\theta(z_1, \dots, z_L)} \\
&= \frac{p_\theta(x_j, \mathbf{z}_{>1}) p_\theta(z_1|x_j, \mathbf{z}_{>1})}{p_\theta(\mathbf{z}_{>1}) p_\theta(z_1|\mathbf{z}_{>1})} \\
&= \dots \\
&= \frac{p_\theta(z_L, x_j)}{p(z_L)} \prod_{l=1}^{L-1} \frac{p_\theta(z_l|x_j, \mathbf{z}_{>l})}{p_\theta(z_l|\mathbf{z}_{>l})} \\
&= p_\theta(x_j) \frac{p_\theta(z_L|x_j)}{p(z_L)} \prod_{l=1}^{L-1} \frac{p_\theta(z_l|x_j, \mathbf{z}_{>l})}{p_\theta(z_l|\mathbf{z}_{>l})}.
\end{aligned} \tag{30}$$

Consequently, combining Eq. (29) and Eq. (30), the true conditional distribution  $p_\theta(\mathbf{z}|\mathbf{x}_r^o)$  can be factorized as:

$$\begin{aligned}
p_\theta(\mathbf{z}|\mathbf{x}_r^o) &\propto p(z_L) \prod_{l=1}^{L-1} p_\theta(z_l|\mathbf{z}_{>l}) \\
&\prod_{\substack{j=1 \\ \text{s.t. } r_j=1}} \left[ \frac{p_\theta(z_L|x_j)}{p(z_L)} \prod_{l=1}^{L-1} \frac{p_\theta(z_l|x_j, \mathbf{z}_{>l})}{p_\theta(z_l|\mathbf{z}_{>l})} \right], \tag{31}
\end{aligned}$$

which can be rewritten as:

$$\begin{aligned}
p_\theta(\mathbf{z}|\mathbf{x}_r^o) &= \left( p(z_L) \prod_{\substack{j=1 \\ \text{s.t. } r_j=1}}^M \frac{p_\theta(z_L|x_j)}{p(z_L)} \right) \\
&\prod_{l=1}^{L-1} \left( p_\theta(z_l|\mathbf{z}_{>l}) \prod_{\substack{j=1 \\ \text{s.t. } r_j=1}}^M \frac{p_\theta(z_l|x_j, \mathbf{z}_{>l})}{p_\theta(z_l|\mathbf{z}_{>l})} \right). \tag{32}
\end{aligned}$$