

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/379411159>

# Mutual Information Guided Diffusion for Zero-Shot Cross-Modality Medical Image Translation

Article in IEEE Transactions on Medical Imaging · March 2024

DOI: 10.1109/TMI.2024.3382043

CITATIONS

11

READS

128

7 authors, including:



Zihao Wang

Harvard University

26 PUBLICATIONS 719 CITATIONS

SEE PROFILE



Tingting Yuan

University of Göttingen

22 PUBLICATIONS 88 CITATIONS

SEE PROFILE



Maxime Sermesant

National Institute for Research in Computer Science and Control

437 PUBLICATIONS 12,348 CITATIONS

SEE PROFILE



Hervé Delingette

National Institute for Research in Computer Science and Control

445 PUBLICATIONS 17,116 CITATIONS

SEE PROFILE

# Zero-shot-Learning Cross-Modality Data Translation Through Mutual Information Guided Stochastic Diffusion

Zihao Wang<sup>1\*</sup>, Yingyu Yang<sup>2</sup>,  
 Maxime Sermesant<sup>2</sup>, Hervé Delingette<sup>2</sup>, Ona Wu<sup>1</sup>  
<sup>1</sup> A. A. Martinos Center for Biomedical Imaging@MGH,  
 Harvard University.  
 149 13th St, Charlestown, MA 02129  
<sup>2</sup> Centre Inria d'Université Côte d'Azur.  
 2004 Rte des Lucioles, 06902 Valbonne, France  
 zwang63@mgh.harvard.edu

February 1, 2023

## Abstract

Cross-modality data translation has attracted great interest in image computing. Deep generative models (*e.g.*, GANs) show performance improvement in tackling those problems. Nevertheless, as a fundamental challenge in image translation, the problem of Zero-shot-Learning Cross-Modality Data Translation with fidelity remains unanswered. This paper proposes a new unsupervised zero-shot-learning method named Mutual Information guided Diffusion cross-modality data translation Model (MIDiffusion), which learns to translate the unseen source data to the target domain. The MIDiffusion leverages a score-matching-based generative model, which learns the prior knowledge in the target domain. We propose a differentiable local-wise-MI-Layer (*LMI*) for conditioning the iterative denoising sampling. The *LMI* captures the identical cross-modality features in the statistical domain for the diffusion guidance; thus, our method does not require retraining when the source domain is changed, as it does not rely on any direct mapping between the source and target domains. This advantage is critical for applying cross-modality data translation methods in practice, as a reasonable amount of source domain dataset is not always available for supervised training. We empirically show the advanced performance of MIDiffusion in comparison with an influential group of generative models, including adversarial-based and other score-matching-based models.

---

\*corresponding: zihao.wang@ieee.org

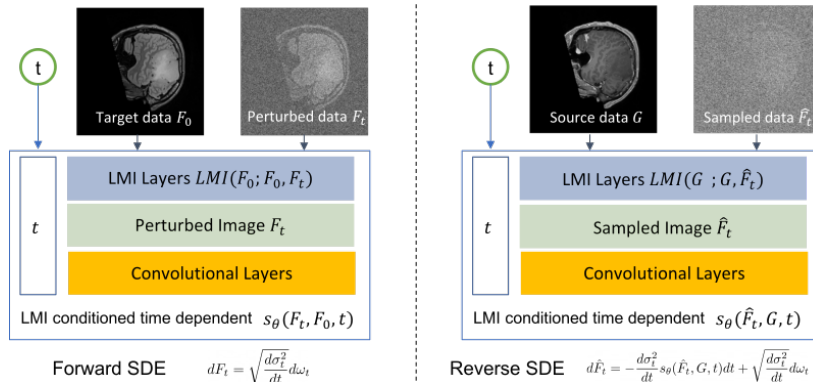


Figure 1: The proposed local-wise mutual information (LMI) guidance for zero-shot stochastic diffusion-based modality translation. Unlike the classical score-matching model, the proposed method conditions on a statistical measure: local-wise mutual information. It measures the LMI between target image  $F$  and its perturbed image  $F_t$  at time  $t$  during training time (left sub-figure) for the forward SDE. At test time (right sub-figure), it calculates the LMI between source image  $G$  and the sampled image  $\hat{F}_t$  at time  $t$  for the reverse SDE.

## 1 Introduction

One often wants to use existing resources to solve new problems without having to develop similar resources for new data modalities to maximize economic utility. For example, software systems to perform a specific task, e.g. brain parcellation, may be designed to work with a particular data modality, e.g. MRI. Ideally, one would be able to translate the new data modality, e.g. CT, into the specific data modality and be able to analyze the new data with the existing software. This method is often termed cross-modality data translation, which has been of particular interest to the generative learning research community in the past few years. Two major types of methods have been proposed: (1) generative adversarial-based translation and (2) mapping-based translation. Those two approaches correspond to the two major machine learning model types: generative and discriminative learning. The mapping-based method directly learns the mapping between different modalities through regression between the source and target modal data points. However, paired datasets of different modalities are not always available in many scenarios. Also, slight pixel-wise misalignments between different modalities may lead to inaccurate translation because mapping-based translation models are designed to build a direct mapping between the source and target modalities. These issues have led to recent research focused on generative adversarial-based translation, which learns the cross-modality mapping through cycle-consistency training, which does not require paired datasets.

Two conceptual metrics have been proposed to evaluate the translation quality:

synthesis realism (see Def. 1) and faithfulness (see Def. 2)[41]. Translation faithfulness refers to how well the translation results maintain the context of the source data intact. Translation realism means how consistent the translation results are with the range of the target modality. Ensuring those two metrics in a Zero-shot-Learning way is challenging as a Zero-shot-trained model does not have any information from the source domain.

Mapping-based methods[60, 16, 26, 18, 43] usually favor translation faithfulness over realism. Many early mapping-based modality translations methods rely on pixel-level modeling between the source and target modality. Generative adversarial network (GAN) [33, 5, 68, 62, 59, 36, 4] based methods were proposed due to various shortcomings of the previous mapping-based approaches. The generative model is usually used for modeling the target modality directly, thus achieving translation realism[1, 69, 8]. This method usually involves complicated adversarial architecture design and modality task-specific loss function design for different translation tasks [41]. Although the generative adversarial-based translation does not need access to paired datasets for training, it still requires source domain data, which may be challenging to collect, leading to insufficient samples to balance the cycle-consistency training. Recent works [15, 14, 31, 6] show that score-based generative models achieve better generation performance than that GAN-based models. Meng *et. al.* [41] proposed SDEdit, which employs a Stochastic Diffusion Model (SDM) [57] to perform image translation that balanced faithfulness and realism in the way of Zero-shot-Learning. Different from conventional GAN or mapping-based models, which are nearly end-to-end generation, an SDM is a score-based generative model [23] that relies on iteratively denoising a diffusion sequence driven by an SDE. Given a datapoint from the source domain, SDEdit perturbs the datapoint with Gaussian noise (perturbation-based guidance). Then a reverse SDE is used by SDEdit to gradually project the disturbed datapoint to the manifold of the target domain. SDEdit has been shown to be superior to GAN-based translation models in terms of parsimony of model structure and complexity of loss function. Although SDEdit overcomes the drawbacks of previous GAN-based translation methods, it remains limited in cross-modality data translation applications. The perturbation-based guidance assumes the source and target domains can be perturbed by noise, which may not be true in specific translation tasks (e.g., microscopy brightfield  $\rightarrow$  darkfield translation). In addition, SDEdit requires optimizing initial time  $t_0$  to find the best interval for perturbation.[41] Muzaffer *et. al.* [46] proposed SynDiff to tackle this issue by introducing a cycle-consistent architecture devised with a bilateral diffusion process. The advantage of SynDiff is that the transformation is semantically consistent, yet the computational cost of computing a bilateral diffusion process is doubled. In fact, although it is claimed to be an unsupervised method, SynDiff needs to pre-train a generator to estimate a paired source image for training purposes, which requires an additional generation module; thus, the model’s overall performance is also influenced by the pre-training quality of the generator.

We propose a new zero-shot unsupervised learning-based cross-modality data translation method based on a stochastic diffusion process. Instead of seeking any

conditioning guidance from the data domain, our model leverages the statistical feature-wise homogeneity for conditioning the diffusion process (see Fig. 1). This allows us to bridge the source and target domain by using their local-wise statistical attributions for cross-modality data translation. The proposed method overcomes the shortcomings of current modality translation approaches:

(1) Unlike GAN-based strategies, the proposed method (MIDiffusion) neither needs inversion optimization in the test step nor requires additional adversarial architectures and loss functions design (cycle consistency, *etc.*).

(2) In contrast to the current SDEdit method, our framework does not require the optimization of the hyper-parameter  $t_0$ , which is a user-selected parameter, to achieve balance between realism and faithfulness.[41]

(3) MIDiffusion is a fully unsupervised Zero-shot-Learning model that can translate unseen modalities into target modalities.

(4) MIDiffusion also empirically shows that the training of an additional generator [46] for diffusion process guidance is not necessary to achieve semantic consistency.

## 2 Related Works

**Generative Learning-based Cross-modality Data Translation** A plethora of works performs cross-modality data translation through conditional generative adversarial networks (cGANs) [72, 71, 58, 54, 77, 4]. Those cGANs-based cross-modality data translation usually requires the generation of a supervised signal for conditioning the generator network (e.g., Zhuge *et. al.* [4] used a register to generate the conditional signal). Some other GANs-based cross-modality data translation methods use cycle consistency training to swap the features between different domains [52, 29, 21, 50].

**Kernel-based Generative Learning** The proposed method falls in the scope of transition-kernel-based generative learning [56, 51, 65, 55, 10, 35]; specifically, it belongs to score-based generative learning [24, 23, 61, 27]. The score-based generative models show comparable data modeling performance to those generative adversarial methods [33, 5, 68, 62, 59, 36, 8, 4, 28].

**Zero-shot Learning** The cross-modality data translation problem without access to the source modal data leads to a Zero-shot-Learning-based cross-modality data translation problem. One challenge of learning-based methods is that their modeling ability is restricted when dealing with unseen data classes [66, 9, 34]. Zero-shot-Learning is a robust learning scheme to deal with such cases when training and test classes are disjoint. One traditional approach of Zero-shot-Learning tries to find a direct projection from image feature space to a semantic space through discriminative methods [48, 2] or generative models [39, 63]. Another popular way is to combine different modalities (such as images, texts, attributes, *etc.*) and learn a non-linear multi-modal embedding [11, 38], to help unseen class recognition. For example, Lin *et. al.* [37] used a GAN-based model to learn a multi-modal consistent semantic representation, and the disentangled domain-invariant features are extracted for unsupervised zero-shot

image-to-image translation.

**Zero-shot Learning-based Cross-modality Data Translation** The GAN-based zero-shot-learning cross-modality data translation models usually modify the latent representation of a pre-trained GAN model (a.k.a. GAN inversion [75, 53, 8, 1]). It needs additional optimization in the testing step, which is computationally expensive and time-consuming.

Perturbation-diffusion-based cross-modality data translation allows Zero-shot learning for cross-modality data translation tasks [13, 23, 32, 57]. The perturbation-diffusion-based methods perform excellently when the numerical features of source and target domains are consistent [41]. Nevertheless, they may fail when the cross-domain appearance features are significantly distinct.

### 3 Theroy

#### 3.1 Cross-modality Data Translation

Mathematically, the cross-modality data translation task between two modalities  $G \in \mathcal{V}$  and  $F \in \mathcal{U}$  can be formalized as:

$$\hat{F} \in \mathcal{U} : \hat{F} = \Phi_F(G) \tag{1}$$

, where  $\Phi_F$  is an operator that maps the data  $G$  in the source domain  $\mathcal{V}$  to the corresponding data  $\hat{F}$ , which is ideally the same as  $F$  in the target domain  $\mathcal{U}$ . Specifically, this generalized form has been broadly applied in image synthesis, stroke painting, image registration, segmentation, *etc.* [49] In this paper, we will mainly focus on 2D image space. However, the proposed method can be extended to 1D or 3D signals without loss of generality.

In zero-shot cross-modality data translation, only samples in target domain  $F \in \mathcal{U}$  are available during the training phase. Given the target samples, the aim of Zero-shot-Learning is to learn  $\Phi_F$  without seeing the  $G$  in training step. Since no source samples are accessible in the training phase, it is important to build and use auxiliary information for domain transfer [64].

#### 3.2 Mutual Information

Mutual Information (MI) measures the dependence of two random variables  $X, Y$ :

$$MI(X, Y) = \iint p(x, y) \log \frac{p(x, y)}{p(x)p(y)} dx dy \tag{2}$$

MI is useful in cross-modality data processing tasks as the statistical features are assumed to be identical. It has been applied to tackle many unsupervised learning problems such as cross-modality data retrieval[25], data representations [22, 74, 17], domain adaptation [42], and cross-modal clustering [40] *etc.*

A particular case of MI is using MI for measuring a random variable itself:  $MI(X, X)$ , which is called *Entropy*.

### 3.3 Score Matching and Its Denoising Equivalent

Different from variational inference-based or likelihood-based training, which attempts to approximate the true probability distribution  $\log p(x)$  of the data, the score-based models learn to represent the distribution  $\log p(x)$  through its partial derivative  $\frac{\partial \log p(x)}{\partial x}$  information (a.k.a score function). The maximization process between a learning model  $s_\theta(x)$  and  $\frac{\partial \log p(x)}{\partial x}$  needs to get an explicit form  $q(x)$  of the accurate distribution  $p(x)$ , which usually remains unknown. Instead of seeking any form  $q(x)$  of the distribution of the target dataset, the denoising score matching method [61] sidesteps the searching of  $q(x)$  by directly estimating the score function through [61, 27]:

$$\arg \min_{\theta} \mathbb{E}_{q_\sigma(x, \hat{x})} \left[ \frac{1}{2} \left\| s_\theta(x) - \frac{\partial \log q_\sigma(\hat{x}|x)}{\partial \hat{x}} \right\|^2 \right] \quad (3)$$

, where  $\frac{\partial \log q_\sigma(\hat{x}|x)}{\partial \hat{x}}$  is the gradient between the clean data  $x$  and a noise-polluted observation  $\hat{x}$ ; as long as the noise is driven by a Gaussian kernel:  $\frac{\partial \log q_\sigma(\hat{x}|x)}{\partial \hat{x}} = \frac{x - \hat{x}}{\sigma^2}$ , the learning model  $s(x)$  aims to learn the process of denoising.[61]

### 3.4 Score-based Generative Modeling with SDEs

Song *et. al.* [57] generalized the above denoising score matching-based models into an SDE framework and unified the generative process by sampling with a reverse SDE [3], which is also a diffusion process. The neural network  $s_\theta$  is trained in an implicit denoising form by treating the noise-adding steps as a diffusion process. The diffusion process driven by a standard normal distribution  $\mathcal{N}(0, 1)$  in a potential  $U_t(x)$  can be modeled by an SDE:

$$dx_t = U_t(x)dt + \sigma_t d\mathbf{w}_t, \quad t \in [0, T] \quad (4)$$

, where  $\sigma_t \in [0, \infty)$  controls the magnitude of the input noise,  $\mathbf{w}_t$  denotes a Wiener process, and  $t \in [0, T]$  is the time start from 0 with infinitesimal incremental to  $t$ . We use a static potential  $U_t(x) = 0$  (a.k.a "Variance Exploding SDE (VE-SDE)" [57]) for modality translation here, which models the magnitude of the data  $x_0$  ( the target modality  $F_0$ ). Adding noise with Eq. 4 is a similar step when using a Gaussian transition kernel to perform score matching in Eq. 3. The training target of a diffusion process defined by the SDE 4 shares a similar training target as Eq. 3 yet with an expectation of time  $t$  sampled uniformly between  $[0, T]$ :  $\mathbb{E}_{t \sim \mathcal{U}(0, T)}$ ; also the whole process becomes a multi-step conditioning  $q(x_t|x_0)$  instead of a single step  $q(\hat{x}|x)$ ,

$$\arg \min_{\theta} \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left\{ \mathbb{E}_{x_0 \sim p_0(x)} \mathbb{E}_{x_t \sim q_{\sigma_t}(x_t, x_0)} \left[ \frac{1}{2} \left\| s_\theta(x) - \frac{\partial \log q_{\sigma_t}(x_t|x_0)}{\partial x_t} \right\|^2 \right] \right\} \quad (5)$$




Figure 2: Example data of different modalities for the three different datasets; from left to right are: *CuRIOUS* dataset of FLAIR imaging, *CuRIOUS* dataset of T1-weighted (T1-w) imaging, *Gold Atlas* dataset of MR imaging, *Gold Atlas* dataset of CT imaging, *IXI* dataset of PD-weighted imaging, *IXI* dataset of T1-w imaging.




Figure 3: Illustration of an example of the proposed operator functioning for 1D functions  $X$  and  $Y$  in the neighborhood  $\delta$  of point  $x_i$ . (I)  $\delta$ -neighborhood of the two functions (upper  $Y$ , lower  $X$ ) at point  $y_i$  and  $x_i$ . (II) functions are processed by the proposed operators defined in Def. 5 and Def. 4. (III) compute the Mutual Information between the paired segments in step (II) using kernel density estimation. (IV) select the maximum value of MI in (III) as  $LMI_\delta(x_i, y_i)$  at point  $x_i$  between  $X$  and  $Y$ .

As long as the matched score model  $s_\theta \sim \frac{\partial \log p(x)}{\partial x}$  learns the distribution score  $\frac{\partial \log p(x)}{\partial x}$ , we can employ a reverse SDE [57, 3, 41] to sample datapoints by inferring a backward in time  $t : T \rightarrow 0$  dynamic process:

$$dx_t = -\frac{d\sigma_t^2}{dt} \left[ \frac{\partial \log p(x)}{\partial x} \right]_t dt + \sqrt{\frac{d\sigma_t^2}{dt}} d\mathbf{w}_t, \quad (6)$$

$$= -\frac{d\sigma_t^2}{dt} s_\theta(x_t, t) dt + \sqrt{\frac{d\sigma_t^2}{dt}} d\mathbf{w}_t, \quad t_{T \rightarrow 0} \in [0, T] \quad (7)$$

, where  $\mathbf{w}_t$  is a Wiener process with infinitesimal negative incremental of time  $dt$ .

## 4 Methods

### 4.1 Diffusion for Cross-modality Data Translation

We can solve the cross-modality data translation problem by adapting the target data  $F$  generation task into the framework of score-matching and then using a perturbed source domain  $G$  to guide (conditioning) the iterative diffusion process [41, 46, 7, 47, 73, 32, 6]. Ideally, we want the generated data  $\hat{F}$  to follow the semantic meaning of the guided data  $G$  and share the features as the ground truth data point  $F$  in the target domain, balancing **realism** and **faithfulness**. [41, 13]

#### 4.1.1 Cross-modality Data Translation Fidelity

**Definition 1** *The generated image  $\hat{F}$  shows **realism** means it is well translated into the target domain  $\mathcal{U}$ :  $\hat{F} \in \mathcal{U}$ .*

**Definition 2** *The generated image  $\hat{F}$  remains **faithfulness** means it is faithfully translated from the guided data:  $\hat{F} \sim G \in \mathcal{V}$ , where  $\sim$  is a similarity measure.*

The relationship between translation realism and faithfulness is related to the domain correlations  $\mathcal{U}$  and  $\mathcal{V}$ . We say a translation achieves high **fidelity** if and only if the translation achieves a balance point between the **realism** and **faithfulness**. However, when the difference in the appearance feature set between the source domain  $G$  and target domain  $F$  becomes too large, a balance point between the two sides may neither be easy to capture nor exist for a satisfying translation. This leads current distribution perturbing-based methods (e.g., SDEdit, etc. [41, 61, 27]) unsuitable for unsupervised cross-modality data translation when the numerical features between the two domains have huge differences.

### 4.2 Mutual Information Guidance in Diffusion Generation

In the Zero-shot-Learning-based translation task, we do not have access to the data in the source domain during the training process. Nevertheless, the local statistical features between the source and target modalities are assumed to be identical. MI maximization has been proven an effective method to empower neural networks to learn non-linear representations [22]. To capture those shared representations and use the extracted information for generation guidance, we propose using MI to measure the local statistical representations in the iterative denoising process.

#### 4.2.1 Local-wise MI

To obtain the semantic information in the data for guidance, we need to convert the original data to statistical representation, as MI is a statistical measure.

Given a datapoint  $X$ , for point  $x_i \in X$  at position  $i$ , the local-wise statistical information at  $i$  can be captured through the probability density function (PDF)  $p_{\delta_{x_i}}(\cdot)$  of the neighborhood area  $\delta_{x_i}$  of  $x_i$ .

Without loss of generality, for the other points  $x_j$  inside the neighborhood area  $\delta_{x_i}$  of  $x_i$ , we can obtain the local-wise statistical information through the PDF  $p_{\delta_{x_j}}(\cdot); j \in \delta_i$ .

**Definition 3** *The local-wise MI (LMI) from data point  $X$  to data point  $Y$  at point  $x_i$  is defined through:*

$$LMI_{\delta}(x_i, y_j) = \sup \iint p_{\delta}(x, y) \log \frac{p_{\delta}(x, y)}{p_{\delta_{x_i}}(x)p_{\delta_{y_j}}(y)} dx dy, \forall y_j \in \delta_{x_i} \quad (8)$$

In the forward steps (training), we can use Eq. 8 as a reference signal to condition each diffusion step; this is achieved by computing the  $LMI(X_0, X_t), t \in [0, T]$  during the training process of the scoring neural network.

**Theorem 1** *The upper bound of the LMI from data point  $X$  to datapoint  $Y$  at location  $i$  is:  $LMI_{\delta}(x_i, y_i) \leq LMI_{\delta}(x_i, x_i)$ , which is the optimum informative match between  $X$  and  $Y$  at point  $x_i$ .*

The Thm. 1 indicates that the  $LMI$  achieves maximum (statistical similarity) when:  $p_{\delta}(X) = p_{\delta}(Y), \forall \delta \in X$ . Thus the  $LMI$  always achieves maximum at the same position between  $X_0$  and  $X_t$  in the training step; Yet, when  $p_{\delta}(X) \neq p_{\delta}(Y)$ , the  $LMI$  achieves the local maximum at position  $j$ , which is located in the neighborhood  $\delta_{y_i}$  of  $y_i$ .

We have defined the  $LMI$  for conditioning the iterative diffusion process. However, the computational cost of  $LMI$  is very high as a statistical measure. The iterative training of the score model requires many steps of discrete time conditioning. It is unrealistic to use a sliding window or patch-wise looping to compute the  $LMI$  between each time steps. To overcome this bottleneck of applying  $LMI$  in score matching models, we developed an efficient method to compute the  $LMI$  in both of the perturbing and denoising steps.

#### 4.2.2 Differentiable local-wise MI Layer

**Definition 4** *Let  $X(i)$  be a function defined in  $\mathbb{R}^n$ , then  $C_{\delta}(X) \in \mathbb{R}^n$  is a **periodic extension** of  $X$  in neighborhood  $K \cdot \delta, K \in \mathbb{Z}$ :*

$$C_{\delta}(X) := \begin{cases} X & \text{for } i \in \delta \\ \text{is periodic of } \delta & \text{for } i \in K \cdot \delta \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

**Definition 5** Let  $Y(i)$  be a function defined in  $\mathbb{R}^n$ , then  $\mathcal{B}_\delta(Y) \in \mathbb{R}^n$  is a  $k$ -steps ( $k \in \mathbb{Z}^n$ ) **sliding extension** of  $Y$  in neighborhood  $k \cdot \delta$ :

$$\mathcal{B}_\delta(Y) := J_\tau Y(i - \tau)$$

$$\text{s.t. } \mathbf{J}_\tau = \begin{cases} 1 & \tau - \frac{\delta}{2} \leq i < \tau + \frac{\delta}{2} \\ 0 & \text{else} \end{cases} \quad (10)$$

for  $\tau : 0 \rightarrow \delta$  with incremental  $\Delta\tau = \delta/k$ .

Definition (5) and (4) can be understood as a 'segment-wise' linear interpolation and nearest-neighbor interpolation [67] of functions  $Y$  and  $X$  in the neighborhood  $\delta$ . A demonstration of the operator applied for 1D functions defined in (Def. 4) and (Def. 5) is shown in Fig. 3 (II), for which the two functions  $X$  and  $Y$  (Fig. 3 (I)) are processed by the operators  $\mathcal{C}$  and  $\mathcal{B}$  respectively.

**Proposition 1** With two given functions  $F \in \mathcal{U}$  and  $G \in \mathcal{U}$  in  $\mathbb{R}^2$ , the LMI (Def. 8) between  $G$  and  $F$  can be computed through the following operator  $\mathcal{S}$ :

$$LMI(G, F) = \max_{\delta} (\mathcal{K}(\mathcal{B}_\delta(F), \mathcal{C}_\delta(G)) \log \frac{\mathcal{K}(\mathcal{B}_\delta(F), \mathcal{C}_\delta(G))}{\mathcal{K}\mathcal{B}_\delta(F)\mathcal{K}\mathcal{C}_\delta(G)})$$

$$\delta \in G, F \quad (11)$$

where  $\mathcal{K}$  is a kernel density estimator, which approximates the PDF.

The Prop. 1 computes the maximum LMI in both of the training and testing steps, which are  $LMI_\delta(F, F)$  and  $LMI_\delta(G, F)$  respectively. The operator given in Prop. 1 can be accelerated by memory copying and parallel reduction in GPGPU [12].

---

**Algorithm 1:** Mutual Information Guided Diffusion

---

**Inputs:**  $G$ ,  $s_\theta$ , time  $t_{T \rightarrow 0} \in [0, T]$ , step size  $dt$

**Output:** Generated high fidelity data  $\hat{F}$

Initialize all parameters, variables;

**while**  $t < T$  **do**

$z \sim \mathcal{N}(0, \mathbf{1})$ ;

$t \leftarrow t + dt$ ;

$\epsilon = \sigma_t^2 - \sigma_{t-dt}^2$ ;

$\hat{F}_{t+1} \leftarrow \hat{F}_t + \epsilon s_\theta(\hat{F}_t, LMI(G; G, \hat{F}_t), t) + \sqrt{\epsilon} z$  ;                      // Euler

    discretization step, guided by  $G$

**end**

---

### 4.3 Embedding the Conditioner into SDE

In the training step (see left sub-figure of Fig. 1), we can condition the noise perturbing process by embedding the defined LMI between a datapoint  $F$  and

Table 1: Quantitative evaluation of the MIDiffusion in comparison with three baselines on three different cross-modality data translation tasks. The best-performing values on the same task being compared are bolded in black italics for highlighting. sup=supervised; unsup=unsupervised

Datasets	Methods	Modalities	SSIM	SSIM	MSE↓	MI↑	PSNR↑	FID↓
			Tar↑	Src↑				
GoldAtlas	CycleGAN[76]	CT→MR	0.04	0.03	614.02	<b>1.16</b>	20.53	<b>202.43</b>
	(sup, few-shot 2%)	MR→CT	0.03	0.02	819.59	1.13	19.08	281.35
	StyleGAN [30]	CT→MR	<b>0.13</b>	0.04	788.76	1.09	20.09	213.47
	(unsup, inversion)	MR→CT	0.08	0.07	570.91	1.12	21.17	<b>170.83</b>
	SDEdit [41]	CT→MR	0.003	0.01	766.40	1.11	19.50	237.27
	(unsup)	MR→CT	0.01	0.04	996.71	1.10	18.58	223.44
	MIDiffusion	CT→MR	0.06	<b>0.11</b>	<b>523.18</b>	1.08	<b>21.66</b>	245.82
(unsup)	MR→CT	<b>0.12</b>	<b>0.08</b>	<b>392.35</b>	<b>1.17</b>	<b>23.03</b>	194.35	
CuRIOUS	CycleGAN	T1→FLAIR	-0.006	0.81	1747.13	<b>1.08</b>	16.04	186.59
	(sup, few-shot ~6%)	FLAIR→T1	0.005	0.02	3145.05	1.05	13.82	331.89
	StyleGAN	T1→FLAIR	0.003	0.12	1880.62	1.04	15.83	261.47
	(unsup, inversion)	FLAIR→T1	-0.003	0.19	1570.83	1.05	16.62	229.73
	SDEdit	T1→FLAIR	0.011	0.01	1558.22	1.04	16.42	<b>131.70</b>
	(unsup)	FLAIR→T1	0.005	0.01	2165.42	1.03	15.14	<b>141.89</b>
	MIDiffusion	T1→FLAIR	<b>0.07</b>	-0.08	<b>1226.40</b>	<b>1.08</b>	<b>17.65</b>	146.77
(unsup)	FLAIR→T1	<b>0.15</b>	<b>0.23</b>	<b>1175.11</b>	<b>1.08</b>	<b>18.02</b>	157.98	
IXI	CycleGAN	PD→T1	<b>0.12</b>	0.14	<b>1154.19</b>	<b>1.17</b>	<b>17.65</b>	141.95
	(sup, few-shot 11%)	T1→PD	0.16	0.16	<b>876.99</b>	<b>1.19</b>	<b>18.86</b>	113.67
	StyleGAN	PD-T1	0.02	0.06	6609.13	1.08	10.17	266.52
	(unsup, inversion)	T1→PD	<b>0.21</b>	<b>0.37</b>	2319.78	<b>1.14</b>	14.65	199.12
	SDEdit	PD-T1	0.09	0.06	1619.14	1.15	16.19	<b>68.6</b>
	(unsup)	T1→PD	0.10	0.06	1753.82	1.16	15.95	<b>80.81</b>
	MIDiffusion	PD-T1	<b>0.11</b>	<b>0.19</b>	<b>1652.81</b>	<b>1.17</b>	16.35	129.12
(unsup)	T1→PD	0.18	<b>0.26</b>	<b>1301.91</b>	1.13	17.13	132.46	

its disturbed data  $F_t$  in to the time-dependent-scoring model  $s_\theta$ , which is guided by  $G$ :

$$\arg \min_{\theta} \mathbb{E}_{t \sim \mathcal{U}(0, T)} \{ \mathbb{E}_{x_0 \sim p_0(x)} \mathbb{E}_{x_t \sim q_{\sigma_t}(x_t, x_0)} \left[ \frac{1}{2} |s_\theta(\hat{F}_t, LMI(F; F, F_t), t) - \frac{\partial \log q_{\sigma_t}(x_t | x_0)}{\partial x_t}|^2 \right] \} \quad (12)$$

In the sampling step (see right sub-figure of Fig. 1), recall that the generative process is an iterative solution of the reverse SDE 6. We can inject the proposed conditioning procedure into the SDE 6:

$$d\hat{F}_t = -\frac{d\sigma_t^2}{dt} s_\theta(\hat{F}_t, LMI(G; G, \hat{F}_t), t) dt + \sqrt{\frac{d\sigma_t^2}{dt}} d\mathbf{w}_t, \quad t_{T \rightarrow 0} \in [0, T] \quad (13)$$

and use an naive Euler-Maruyama numerical solver (see Algorithm 1) to solve the SDE.

## 5 Experiments and Evaluation

We show the following features of the proposed method in this section: (1) The proposed zero-shot unsupervised learning method outperforms GAN-based cross-modality data translation models in both the zero-shot and few-shot supervised tracks. (2) Our method achieves high cross-modality data translation fidelity compared with the state-of-the-art diffusion method. (3) The proposed *LMI*-guided diffusion model enables cross-modality data translation semantic consistency in the practical application.

### 5.1 Evaluation Datasets

We introduce three public datasets that cover different cross-modality data translation tasks. Fig. 2 shows six samples collected from three different datasets: (1) The *Gold Atlas dataset* includes CT and T1-weighted (T1w) and T2-weighted (T2w) magnetic resonance imaging (MRI) data from 19 male patients. The pelvic area was the subject of the imaging. All the CT images had been deformably registered to the corresponding MRI [45]. The datasets were collected from three different sites [44]. We resampled all the datasets using the SimpleITK package to a uniform voxel size ( $0.875mm \times 0.875mm \times 3mm$ ), and selected slices every  $15mm$  interval to build the training set (993 CT-MR slice pairs from 15 patients) and the testing set (227 CT-MR slice pairs from 4 patients).

(2) The publicly available *CuRIOUS dataset* consists of 22 subjects with low-grade glioma [70]. The original dataset includes T1w and FLAIR MRI scan pairs, as well as unregistered ultrasound scans. All the images were collected during routine clinical exams. The original scans had been resampled to  $256 \times 256 \times 288$  voxels at an isotropic voxel size of  $0.5mm^3$  [19]. We chose the T1w and FLAIR MRI scan pairs to build our cross-modality image translation task. All the volume pairs were resampled to  $128 \times 128 \times 288$  voxels with a spacing size of  $1.0mm \times 1.0mm \times 0.5mm$  by the SimpleITK package. As a result, 1168 FLAIR-T1w slice pairs were picked in every  $5mm$  from the original voxels to generate the training (952 pairs from 17 subjects) and testing (216 pairs from 5 subjects) datasets.

(3) *IXI dataset*<sup>1</sup> includes pre-aligned 600 images from normal subjects. The full dataset includes five different modalities: T1w, T2w, PD-weighted, magnetic resonance angiograms (MRA), and Diffusion-weighted MRI. In our experiment, we perform PD-T1w modal translation tasks. A training set of 300 slices (sampled from 100 subjects) and a testing set of 75 slices (sampled from 25 subjects) were generated using a subset of the *IXI dataset*.

### 5.2 Baselines

We select two representative GAN-based image translation and synthesis approaches and one state-of-the-art diffusion-based method as baselines for comparison. We first compare our zero-shot-Learning-based method with a few-shot

<sup>1</sup><https://brain-development.org/ixi-dataset/>




Figure 4: Qualitative performance of different methods on CuRIOUS, GoldAtlas, and IXI datasets. The top row shows the original images. The second to the last rows show: CycleGAN (few-shot sup), StyleGAN(unsup zero-shot with GAN inversion), SDEdit(zero-shot unsup learning with SDE-based perturbing guidance), and the proposed MIDiffusion (unsup zero-shot with  $LMI$  guidance).

learning-based CycleGAN translation model [76]. The CycleGAN in this experiment will be allowed to see both the full target domain dataset and a small group (about 2% for *Gold Atlas dataset*, 6% for *CuRIOUS dataset* and 11% for *IXI dataset*) of the source domain dataset. However, our proposed method will see only the data in the target domain.

The second baseline method is a GAN inversion-based approach. A StyleGAN2-ADA [30] is allowed to see the target domain training data. The out-domain guided generation is performed through 5000 steps of optimization of inversion in the latent space of the trained StyleGAN2-ADA [75]. An extra encoder network is introduced in the generation step for inversion.

The third baseline model is SDEdit [41], which works for an ablation study. The SDEdit is also a diffusion model-based translation method but uses a distribution perturbation guidance. Whenever possible, we use a default training setup for all the baselines provided in the papers. The experiments were implemented based on publicly available open-source code in the same experiment environment.

### 5.3 Quantitative Performance

We evaluate the different methods based on five metrics, including three image quality measures: SSIM (structural similarity index measure), PSNR (peak signal-to-noise ratio), MSE (mean square error), and two statistical similarity measures: FID (Fréchet Inception Distance) [20] and MI (mutual information). To study translation faithfulness (Def. 2) of each method, we report the SSIM between the translation results  $\hat{F}$  and the guidance data  $G$  (SSIM-Src) and the target data  $F$  (SSIM-Tar). The FID and MI are used to evaluate the translation

similarity (Def. 1) statistically between  $\hat{F}$  and  $F$  for different methods. The FID evaluates the generation similarity in the feature level between the  $\hat{F}$  and  $F$ . The MI assesses the statistical similarity between the paired  $\hat{F}$  and  $F$  images. We show the performance based on those metrics of the four methods in Table. 1.

### 5.3.1 Comparison with few-shot training

We compare our method with the supervised few-shot trained CycleGAN (source domain  $F \in \mathcal{U}$  is visible) on each dataset. We notice that the size of the source datasets for the few-shot training differs between different datasets. Regarding the translation errors (measured by SSIM, MSE, and PSNR), the proposed zero-shot trained MIDiffusion model outperforms the few-shot trained CycleGAN model on both the GoldAtlas (2% of training data from source, 100% from target) and CuRIOUS (6% of training data from source, 100% from target) datasets. However, this advantage of the MIDiffusion vanishes on the IXI dataset for the MSE and PSNR metrics. This can be attributed to the 11% training data from the source domain being sufficient for the CycleGAN to surpass the performance of the zero-shot-learning models.

### 5.3.2 Comparison with zero-shot training

As shown in Table 1, the proposed method performs better than the other one-shot training-based methods in terms of the SSIM-Src, SSIM-Tar, and MSE, which means the generated data of MIDiffusion achieves higher semantic faithfulness with regard to both the source and target domain. Our model also outperforms the other methods in terms of the PSNR and MI. This implies our method shows translation realism 1 with respect to the target data in pairwise comparison. In terms of similarity between the generated dataset and the target dataset, we see that the SDEdit has an overall lower FID score. Yet, in considering our generation target of fidelity translation, the SDEdit fails in faithfulness translation. Thus the translation results of SDEdit performs worse. In contrast, the proposed method achieves good realism (lowest FID score apart from SDEdit) while keeping the semantic meaning from the guidance to the target (higher SSIM, lower MSE). The cross-modality data translation method with the highest fidelity is therefore the MIDiffusion.

## 5.4 Qualitative Performance

Fig. 4 shows translated results from different methods on the test dataset as well as the groundtruth (top rows). The first two columns in each dataset group represent different modalities, and the last two columns correspond to their zoomed-in details. Overall, the proposed method achieves the best translation fidelity among the compared methods. The anatomical structures are more faithfully represented in the MIDiffusion translation results compared to the results from the baseline models. The SDEdit and StyleGAN fail to translate

the source images with identical features. In addition, the few-shot trained CycleGAN clearly underperforms the zero-shot trained MIDiffusion model when trained with insufficient source data on the CuRIOUS (2%) and GoldAtlas (6%) dataset. This observation is consistent with the quantitative results shown in Table 1. Overall, the MIDiffusion method outperforms the baseline methods in both anatomical consistency and appearance similarity.

## 6 Conclusion

This paper presents a novel local-wise mutual information-guided diffusion model named MIDiffusion for cross-modality data translation. Different from current cycle-consistency training, the MIDiffusion does not require seeing the source dataset for training. Unlike GAN inversion methods that require iterative optimization during the generation, the MIDiffusion does not need online optimization in the test steps. Our method introduces a new conditioner that achieves high-fidelity generation but does not need any extra training on the main diffusion flow. This feature is different from supervised guidance [15], which needs to train an extra model to pose conditional guidance on the diffusion flow. The proposed diffusion model is more robust than another diffusion model (SDEdit) in terms of translation faithfulness, thanks to the *LMI* guidance signal of MIDiffusion.

**Limitations** However, the MIDiffusion requires hundreds of times iterative solutions of SDE, which takes dozens of seconds for translating a single image on a GPU. Effectively reducing the number of sampling steps without compromising the translation fidelity would be meaningful for future work.

## References

- [1] R. Abdal, Y. Qin, and P. Wonka. Image2stylegan++: How to edit the embedded images? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [2] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid. Label-embedding for image classification. *IEEE transactions on pattern analysis and machine intelligence*, 38(7):1425–1438, 2015.
- [3] B. D. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [4] M. Arar, Y. Ginger, D. Danon, A. H. Bermano, and D. Cohen-Or. Un-supervised multi-modal image registration via geometry preserving image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [5] K. Armanious, C. Jiang, M. Fischer, T. Küstner, T. Hepp, K. Nikolaou,

- S. Gatidis, and B. Yang. Medgan: Medical image translation using gans. *Computerized medical imaging and graphics*, 79:101684, 2020.
- [6] G. Batzolis, J. Stanczuk, C.-B. Schönlieb, and C. Etmann. Non-uniform diffusion models, 2022.
- [7] G. Batzolis, J. Stanczuk, C.-B. Schönlieb, and C. Etmann. Non-uniform diffusion models, 2022.
- [8] A. Brock, T. Lim, J. Ritchie, and N. Weston. Neural photo editing with introspective adversarial networks. In *International Conference on Learning Representations*, 2017.
- [9] M. Bucher, S. Herbin, and F. Jurie. Generating visual representations for zero-shot classification. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 2666–2673, 2017.
- [10] A. L. Caterini, A. Doucet, and D. Sejdinovic. Hamiltonian variational auto-encoder, 2018.
- [11] U. Chaudhuri, B. Banerjee, A. Bhattacharya, and M. Datcu. A simplified framework for zero-shot cross-modal sketch data retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 182–183, 2020.
- [12] J. Cheng, M. Grossman, and T. McKercher. *Professional CUDA c programming*. John Wiley & Sons, 2014.
- [13] S.-I. Cheng, Y.-J. Chen, W.-C. Chiu, H.-Y. Tseng, and H.-Y. Lee. Adaptively-realistic image generation from stroke and sketch with diffusion model, 2022.
- [14] J. Choi, S. Kim, Y. Jeong, Y. Gwon, and S. Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models, 2021.
- [15] P. Dhariwal and A. Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794. Curran Associates, Inc., 2021.
- [16] Z. Fu and W. Zhang. Research on image translation between sar and optical imagery. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, I-7, 07 2012.
- [17] Y. Guo, B. Liu, and D. Zhao. Online continual learning through mutual information maximization. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 8109–8126. PMLR, 17–23 Jul 2022.

- [18] X. Han. Mr-based synthetic ct generation using a deep convolutional neural network method. *Medical physics*, 44(4):1408–1419, 2017.
- [19] A. Hering, L. Hansen, T. C. W. Mok, A. C. S. Chung, H. Siebert, S. Häger, A. Lange, S. Kuckertz, S. Heldmann, W. Shao, S. Vesal, M. Rusu, G. Sonn, T. Estienne, M. Vakalopoulou, L. Han, Y. Huang, P.-T. Yap, M. Brudfors, Y. Balbastre, S. Joutard, M. Modat, G. Lifshitz, D. Raviv, J. Lv, Q. Li, V. Jaouen, D. Visvikis, C. Fourcade, M. Rubeaux, W. Pan, Z. Xu, B. Jian, F. De Benetti, M. Wodzinski, N. Gunnarsson, J. Sjölund, D. Grzech, H. Qiu, Z. Li, A. Thorley, J. Duan, C. Großbröhmer, A. Hoopes, I. Reinertsen, Y. Xiao, B. Landman, Y. Huo, K. Murphy, N. Lessmann, B. van Ginneken, A. V. Dalca, and M. P. Heinrich. Learn2reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning, 2021.
- [20] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [21] Y. Hiasa, Y. Otake, M. Takao, T. Matsuoka, K. Takashima, A. Carass, J. L. Prince, N. Sugano, and Y. Sato. Cross-modality image synthesis from unpaired data using cyclegan. In *International workshop on simulation and synthesis in medical imaging*, pages 31–41. Springer, 2018.
- [22] R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2019.
- [23] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020.
- [24] J. Ho and T. Salimans. Classifier-free diffusion guidance, 2022.
- [25] T. Hoang, T.-T. Do, T. V. Nguyen, and N.-M. Cheung. Multimodal mutual information maximization: A novel approach for unsupervised deep cross-modal hashing. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–14, 2022.
- [26] T. Huynh, Y. Gao, J. Kang, L. Wang, P. Zhang, J. Lian, and D. Shen. Estimating ct image from mri data using structured random forest and auto-context model. *IEEE Transactions on Medical Imaging*, 35(1):174–183, 2016.

- [27] A. Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005.
- [28] L. Jiang, M. Xu, X. Wang, and L. Sigal. Saliency-guided image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16509–16518, June 2021.
- [29] S. Joshi, R. Osuala, C. Martín-Isla, V. M. Campello, C. Sendra-Balcells, K. Lekadir, and S. Escalera. nn-unet training on cyclegan-translated images for cross-modal domain adaptation in biomedical imaging. In *International MICCAI Brainlesion Workshop*, pages 540–551. Springer, 2022.
- [30] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. Analyzing and improving the image quality of stylegan, 2019.
- [31] B. Kawar, M. Elad, S. Ermon, and J. Song. Denoising diffusion restoration models, 2022.
- [32] B. Kawar, M. Elad, S. Ermon, and J. Song. Denoising diffusion restoration models. *arXiv preprint arXiv:2201.11793*, 2022.
- [33] V. Kearney, B. P. Ziemer, A. Perry, T. Wang, J. W. Chan, L. Ma, O. Morin, S. S. Yom, and T. D. Solberg. Attention-aware discrimination for mr-to-ct image translation using cycle-consistent generative adversarial networks. *Radiology. Artificial Intelligence*, 2(2), 2020.
- [34] H. C. Kuchibhotla, S. S. Malagi, S. Chandhok, and V. N. Balasubramanian. Unseen classes at a later time? no problem, 2022.
- [35] D. Levy, M. D. Hoffman, and J. Sohl-Dickstein. Generalizing hamiltonian monte carlo with neural networks, 2017.
- [36] X. Li, Z. Du, Y. Huang, and Z. Tan. A deep translation (gan) based change detection network for optical and sar remote sensing images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 179:14–34, 2021.
- [37] J. Lin, Y. Xia, S. Liu, S. Zhao, and Z. Chen. Zstgan: An adversarial approach for unsupervised zero-shot image-to-image translation. *Neurocomputing*, 461:327–335, 2021.
- [38] K. Lin, X. Xu, L. Gao, Z. Wang, and H. T. Shen. Learning cross-aligned latent embeddings for zero-shot cross-modal retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11515–11522, 2020.
- [39] Y. Long, L. Liu, F. Shen, L. Shao, and X. Li. Zero-shot learning using synthesised unseen visual data with diffusion regularisation. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2498–2512, 2017.

- [40] Y. Mao, X. Yan, Q. Guo, and Y. Ye. Deep mutual information maximin for cross-modal clustering. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(10):8893–8901, May 2021.
- [41] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2022.
- [42] L. Meng, H. Su, C. Lou, and J. Li. Cross-domain mutual information adversarial maximization. *Engineering Applications of Artificial Intelligence*, 110:104665, 2022.
- [43] D. Nie, X. Cao, Y. Gao, L. Wang, and D. Shen. Estimating ct image from mri data using 3d fully convolutional networks. In *Deep Learning and Data Labeling for Medical Applications*, pages 170–178. Springer, 2016.
- [44] T. Nyholm, S. Svensson, S. Andersson, J. Jonsson, M. Sohlén, C. Gustafsson, E. Kjellén, K. Söderström, P. Albertsson, L. Blomqvist, B. Zackrisson, L. E. Olsson, and A. Gunnlaugsson. MR and CT data with multiobserver delineations of organs in the pelvic area-part of the gold atlas project. *Med. Phys.*, 45(3):1295–1300, Mar. 2018.
- [45] T. Nyholm, S. Svensson, S. Andersson, J. Jonsson, M. Sohlén, C. Gustafsson, E. Kjellen, P. Albertsson, L. Blomqvist, B. Zackrisson, L. E. Olsson, and A. Gunnlaugsson. Gold atlas - male pelvis - gentle radiotherapy, May 2017.
- [46] M. Ozbey, S. U. Dar, H. A. Bedel, O. Dalmaz, S. Ozturk, A. Gungor, and T. Cukur. Unsupervised medical image translation with adversarial diffusion models, 2022.
- [47] O. Özdenizci and R. Legenstein. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *arXiv preprint arXiv:2207.14626*, 2022.
- [48] M. Palatucci, D. Pomerleau, G. E. Hinton, and T. M. Mitchell. Zero-shot learning with semantic output codes. *Advances in neural information processing systems*, 22, 2009.
- [49] Y. Pang, J. Lin, T. Qin, and Z. Chen. Image-to-image translation: Methods and applications, 2021.
- [50] O. Patashnik, D. Danon, H. Zhang, and D. Cohen-Or. Balagan: Cross-modal image translation between imbalanced domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2659–2667, 2021.
- [51] T. Salimans, D. P. Kingma, and M. Welling. Markov chain monte carlo and variational inference: Bridging the gap. In *ICML*, pages 1218–1226, 2015.

- [52] A. Sharma and N. Jindal. Cross-modality breast image translation with improved resolution using generative adversarial networks. *Wireless Personal Communications*, 119(4):2877–2891, Mar. 2021.
- [53] Y. Shi, X. Yang, Y. Wan, and X. Shen. Semanticstylegan: Learning compositional generative priors for controllable image synthesis and editing, 2021.
- [54] A. Sikka, Skand, J. S. Virk, and D. R. Bathula. Mri to pet cross-modality translation using globally and locally aware gan (gla-gan) for multi-modal diagnosis of alzheimer’s disease, 2021.
- [55] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models, 2020.
- [56] J. Song, S. Zhao, and S. Ermon. A-nice-mc: Adversarial training for mcmc, 06 2017.
- [57] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [58] R. Toda, A. Teramoto, M. Kondo, K. Imaizumi, K. Saito, and H. Fujita. Lung cancer CT image generation from a free-form sketch using style-based pix2pix for data augmentation. *Scientific Reports*, 12(1), July 2022.
- [59] H. Toriya, A. Dewan, and I. Kitahara. Sar2opt: Image alignment between multi-modal images using generative adversarial networks. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 923–926. IEEE, 2019.
- [60] P. A. Van Den Elsen, E.-J. D. Pol, T. S. Sumanaweera, P. F. Hemler, S. Napel, and J. R. Adler. Grey value correlation techniques used for automatic matching of ct and mr brain and spine images. In *Visualization in Biomedical Computing 1994*, volume 2359, pages 227–237. SPIE, 1994.
- [61] P. Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- [62] L. Wang, B. Goldluecke, and C. Anklam. L2r gan: Lidar-to-radar translation. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [63] W. Wang, Y. Pu, V. Verma, K. Fan, Y. Zhang, C. Chen, P. Rai, and L. Carin. Zero-shot learning via class-conditioned deep generative models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [64] W. Wang, V. W. Zheng, H. Yu, and C. Miao. A survey of zero-shot learning: Settings, methods, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–37, 2019.

- [65] Z. Wang and H. Delingette. Quasi-symplectic langevin variational autoencoder, 2020.
- [66] Z. Wang, C. Vandersteen, C. Raffaelli, N. Guevara, F. Patou, and H. Delingette. One-shot learning for landmarks detection. In S. Engelhardt, I. Oksuz, D. Zhu, Y. Yuan, A. Mukhopadhyay, N. Heller, S. X. Huang, H. Nguyen, R. Sznitman, and Y. Xue, editors, *Deep Generative Models, and Data Augmentation, Labelling, and Imperfections*, pages 163–172, Cham, 2021. Springer International Publishing.
- [67] G. Wolberg. *Digital image warping*. Los Alamitos, Calif : IEEE Computer Society Press, US, 1990.
- [68] J. M. Wolterink, A. M. Dinkla, M. H. Savenije, P. R. Seevinck, C. A. van den Berg, and I. Išgum. Deep mr to ct synthesis using unpaired data. In *International workshop on simulation and synthesis in medical imaging*, pages 14–23. Springer, 2017.
- [69] Z. Wu, D. Lischinski, and E. Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12863–12872, June 2021.
- [70] Y. Xiao, H. Rivaz, M. Chabanas, M. Fortin, I. Machado, Y. Ou, M. P. Heinrich, J. A. Schnabel, X. Zhong, A. Maier, W. Wein, R. Shams, S. Kadoury, D. Drobny, M. Modat, and I. Reinertsen. Evaluation of mri to ultrasound registration methods for brain shift correction: The curious2018 challenge. *IEEE Transactions on Medical Imaging*, 39(3):777–786, 2020.
- [71] Q. Yang, N. Li, Z. Zhao, X. Fan, E. I.-C. Chang, and Y. Xu. Mri cross-modality neuroimage-to-neuroimage translation, 2018.
- [72] Q. Yang, N. Li, Z. Zhao, X. Fan, E. I.-C. Chang, and Y. Xu. MRI cross-modality image-to-image translation. *Scientific Reports*, 10(1), Feb. 2020.
- [73] M. Zhao, F. Bao, C. Li, and J. Zhu. Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations, 2022.
- [74] X. Zheng, X. Fei, L. Zhang, C. Wu, F. Chao, J. Liu, W. Zeng, Y. Tian, and R. Ji. Neural architecture search with representation mutual information. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11912–11921, June 2022.
- [75] J. Zhu, Y. Shen, D. Zhao, and B. Zhou. In-domain gan inversion for real image editing. In A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, editors, *Computer Vision – ECCV 2020*, pages 592–608, Cham, 2020. Springer International Publishing.

- [76] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.
- [77] H. Zhuge, B. Summa, J. Hamm, and J. Q. Brown. Deep learning 2d and 3d optical sectioning microscopy using cross-modality pix2pix cgan image translation. *Biomed. Opt. Express*, 12(12):7526–7543, Dec 2021.