

Data-Efficient Unsupervised Interpolation Without Any Intermediate Frame for 4D Medical Images

JungEun Kim^{1*} Hangyul Yoon^{1*} Geondo Park¹ Kyungsu Kim^{2†} Eunho Yang^{1,3}

¹Korea Advanced Institute of Science and Technology (KAIST)

²Massachusetts General Hospital and Harvard Medical School ³AITRICS

{jungeun122333, hangyulmd, geondopark, eunhoym}@kaist.ac.kr kskim.doc@gmail.com

Abstract

4D medical images, which represent 3D images with temporal information, are crucial in clinical practice for capturing dynamic changes and monitoring long-term disease progression. However, acquiring 4D medical images poses challenges due to factors such as radiation exposure and imaging duration, necessitating a balance between achieving high temporal resolution and minimizing adverse effects. Given these circumstances, not only is data acquisition challenging, but increasing the frame rate for each dataset also proves difficult. To address this challenge, this paper proposes a simple yet effective Unsupervised Volumetric Interpolation framework, UVI-Net. This framework facilitates temporal interpolation without the need for any intermediate frames, distinguishing it from the majority of other existing unsupervised methods. Experiments on benchmark datasets demonstrate significant improvements across diverse evaluation metrics compared to unsupervised and supervised baselines. Remarkably, our approach achieves this superior performance even when trained with a dataset as small as one, highlighting its exceptional robustness and efficiency in scenarios with sparse supervision. This positions UVI-Net as a compelling alternative for 4D medical imaging, particularly in settings where data availability is limited. The code is available at [UVI-Net](#).

1. Introduction

Video Frame Interpolation (VFI) has been a cornerstone in the realm of video processing, enriching motion visualization by generating intermediate frames. This method primarily relies on intermediate frame supervision, where known frames are used as references to create new intermediate frames. However, applying these VFI methods to 4D medical imaging is not trivial. While the principles of frame interpolation hold the potential for enhancing med-

ical diagnostics and treatments [5, 21, 23, 41, 42, 66, 70], the unique constraints and requirements of medical imaging present challenges.

One significant challenge lies in obtaining a sufficient dataset. Unlike general domain videos, 4D medical images are captured for specific clinical purposes from a relatively small pool of individuals. Similarly, acquiring intermediate frames per image is also hampered by limitations and risks associated with medical imaging modalities.

Computed tomography (CT) exposes patients to elevated radiation levels, potentially increasing the risk of secondary cancer [67]. Similarly, magnetic resonance imaging (MRI) faces the obstacle of lengthy scan times, lasting up to an hour [56], presenting both logistical challenges and issues related to patient comfort. Furthermore, the quality of ground truth intermediate frames in medical imaging is often compromised due to factors such as patient movement, unstable breathing, and the difficulty of maintaining a stable position during prolonged scans [7, 45], limiting data variety and accessibility for research.

In light of these challenges, we present the following question: “Can a VFI model be trained without depending on *any* ground truth intermediate frames?”. Unlike other previous unsupervised approaches in the 2D general domain [36, 39, 54] that interpolate frames given the multiple frame sequences, we address the task of freely interpolating between two given frames *without any intermediate frames*. To achieve this, we propose a straightforward yet effective framework to VFI in medical imaging. By interpolating the flow between two frames with a two-stage process and cycle-consistency constraint, our framework can effectively operate even with a video composited with two frames (i.e., only images of the start and end points exist), entirely in an unsupervised manner. In the initial stage, virtual samples are generated from the two real input images. Subsequently, the real images are reconstructed based on these virtual intermediate samples. This reconstruction process incorporates the candidate images and warped context-

*Equal Contribution †Correspondence to

tual information in multiple scales from the input images. Through this cyclic interpolation approach, we successfully minimize discrepancies between the generated and the actual images by using the real images as a form of pseudo-supervision.

Our proposed method has achieved state-of-the-art results in unsupervised VFI for 4D medical imaging, outperforming the existing techniques with a substantial gap. Our approach also consistently outperforms even for existing supervised methods. Remarkably, our model shows competitive or even superior performance when trained with a minimal training dataset size of just one, contrasting with other baselines that require full datasets, typically exceeding 60 in size. Additionally, the unsupervised nature of our model allows for further performance enhancements through instance-specific optimization. This process involves briefly fine-tuning the model using each test sample during the inference stage, potentially yielding even more refined results.

In summary, our contributions are three-fold:

- We introduce a simple yet effective unsupervised VFI approach for 4D medical imaging. Our methodology leverages cycle consistency constraints within the temporal dimension, thereby obviating the need for ground truth data typically required for interpolated images.
- Our approach achieves state-of-the-art performance, surpassing other unsupervised and supervised interpolation methods. This is accomplished without the instance-specific optimization, which could be employed as a viable option to enhance performance.
- The robustness of our model is particularly evident under conditions of limited data availability, as demonstrated by the increasing performance margin relative to other methods when the dataset size is reduced.

2. Related Works

2.1. Video interpolation

Many studies in the field of video interpolation have been conducted, with a significant emphasis on frame rate up-sampling for natural scene videos [43, 50, 53]. These studies typically rely on ground truth intermediate frames for training [11, 25, 26, 49, 51, 71, 75, 78]. While some studies have explored alternative approaches that do not rely on ground truth intermediate frames, they involve synthesizing frames between a given sequence of intermediate frames [36, 39, 54] or utilize the information from specialized devices, such as event camera [19]. Consequently, applying these methods in settings like our study presents a challenge, as there are no intermediate frames available for synthesis. Furthermore, validation of these methods is restricted to 2D frames, and they encounter challenges when directly applied to volume sequences. This is primarily at-

Data Type	Name	# of Total Inter
2D Natural	UCF101 [64]	2,374,290
	X4K1000FPS [59]	277,704
	Adobe240-fps [65]	79,768
	Vimeo90K [72]	73,171
	ATD-12K [61]	12,000
3D Medical	ACDC [6]	2,556
	4D-Lung [22]	648

Table 1. Comparison of representative 2D VFI datasets with 3D medical VFI datasets in our study. The last column indicates the total number of intermediate frames, representing the sum of intermediate frame counts across each dataset.

tributable to the markedly lower availability of intermediate frames within such datasets, as elucidated in Tab. 1.

Medical 4D image interpolation. To address the above challenges, frame interpolation methods specifically focused on 4D medical images are driven. Several recent works [16, 17] have attempted to interpolate medical 4D images, but these methods rely on the availability of ground-truth intermediate images for training. Although Kim and Ye [30] proposed an interpolation approach without using the authentic intermediate frames, they do not incorporate an unsupervised learning technique for the interpolated samples. Instead, their method involves a post-hoc multiplication of the flow calculation model, which is prone to spatial distortion. This weakness arises since the underlying network does not account for the structural smoothness between two samples during network training. As a result, the scaled calculated flow fails to capture the spatial continuity of intermediate samples beyond the samples provided by authentic frames. Furthermore, since the method focus solely on warping without incorporating image synthesis, they encounter specific issues if a voxel is displaced to a new location without replacement at the original site. Specifically, it results in the voxel appearing twice in the backward-warped frame [37, 40], or a hole at the original location in the forward-warped frame [48]. To overcome the limitation of nonexistent training for intermediate images and warping procedure, we propose a novel network incorporating pseudo-supervision, including an image synthesis network to ensure the integrity of intermediate images.

2.2. Learning optical flow

Optical flow learning is crucial in the video and medical domain. Various learning methods have been extensively investigated [8, 63, 74] aiming to estimate optical flow. However, they require a ground truth optical flow for training, which is limited in availability. To address this limitation, some methods [2–4, 24, 27, 28, 31, 32, 35, 38, 69] have been developed to compute the similarity between the warped image and a fixed reference to train networks, allowing training without ground truth optical flow.

3. Background

We first briefly introduce the necessary background on the flow calculation model in Sec. 3.1 and the existing unsupervised interpolation approaches in Sec. 3.2.

3.1. Flow calculation model

Suppose we are given two input images I_0 and I_1 at time $T = 0$ and $T = 1$, respectively. Our main objective is to predict the intermediate image \hat{I}_t at time $T = t$ within the range of 0 to 1, given I_0 and I_1 , without explicit supervision. An intuitive approach is to train a neural network to directly generate voxel values of \hat{I}_t without explicitly computing coordinate transformation. However, the generation models such as generative adversarial networks (GAN) [12, 46, 47, 73] typically require a large amount of training data, making them impractical for the medical domain where data is limited. In contrast, flow calculation models [25, 54, 71] can generate 3D images using only two real input images. Given these advantages, we employ flow-based methods for this task, as they are capable of generating 3D images using only two real input images.

Flow-based interpolation approaches employ a flow calculation model \mathcal{F}^θ with model parameters θ to obtain a coordinate transformation map between two target samples. Given I_0 and I_1 , the flow calculation model \mathcal{F}^θ takes I_0 and I_1 as sequential inputs and provides a coordinate transformation map $\phi_{0 \rightarrow 1}^\theta$. The objective of the flow calculation model \mathcal{F}^θ is to warp I_0 into $\hat{I}_{0 \rightarrow 1} := I_0 \circ \phi_{0 \rightarrow 1}^\theta$ such that it matches I_1 , where \circ indicates spatial transformation.

To train flow calculation models, a warping loss $\mathcal{L}_{warp}(I_0, I_1)$ is used, which ensures the quality of computed optical flow. The warping loss is defined based on the warped images $I_1 \circ \phi_{1 \rightarrow 0}^\theta$ and $I_0 \circ \phi_{0 \rightarrow 1}^\theta$, which corresponds to I_0 and I_1 , respectively. \mathcal{L}_{warp} can be expressed as

$$\begin{aligned} \mathcal{L}_{warp}^\theta(I_0, I_1) = & \mathcal{L}_{smth}(\phi_{0 \rightarrow 1}^\theta) + \mathcal{L}_{image}(I_1, I_0 \circ \phi_{0 \rightarrow 1}^\theta) \\ & + \mathcal{L}_{smth}(\phi_{1 \rightarrow 0}^\theta) + \mathcal{L}_{image}(I_0, I_1 \circ \phi_{1 \rightarrow 0}^\theta), \end{aligned} \quad (1)$$

where \mathcal{L}_{smth} is a smoothness term that promotes similar flow values among neighboring voxels, and \mathcal{L}_{image} ensures alignment between two images. Typically, we utilize the sum of normalized cross-correlation (NCC) [3] and Charbonnier [9] losses as the \mathcal{L}_{image} , since NCC has extensively used in 3D medical flow calculation works [3, 58, 77], and Charbonnier loss is a common choice in previous VFI works [26, 33, 51, 58, 68, 77]. The losses are defined as:

$$\mathcal{L}_{smth}(\phi) = \|\nabla\phi\|_2 \quad (2)$$

$$\mathcal{L}_{image}(I, \hat{I}) = -NCC(I, \hat{I}) + \sqrt{(I - \hat{I})^2 + \epsilon^2}, \quad (3)$$

where $\nabla\phi$ denotes the flow gradient, and ϵ represents a small constant.

The fully learned flow calculation model, denoted as $\phi_{0 \rightarrow 1}^{\theta^*}$, is earned by minimizing the warping loss $\mathcal{L}_{warp}^\theta(I_0, I_1)$ with respect to the model parameter θ . The calculated flow can be formulated as:

$$\phi_{0 \rightarrow 1}^{\theta^*} \quad \text{s.t.} \quad \theta^* := \arg \min_{\theta} \sum_{(I_0, I_1) \in \mathcal{D}} \mathcal{L}_{warp}^\theta(I_0, I_1), \quad (4)$$

where \mathcal{D} indicates the training set containing the pairs of I_0 and I_1 .

3.2. Previous unsupervised VFI approaches

Methodology. If the flow from I_0 to the intermediate target sample I_t can be ideally acquired as $\phi_{0 \rightarrow t}$, the corresponding \hat{I}_t can also be obtained (i.e., $\hat{I}_t = I_0 \circ \phi_{0 \rightarrow t}$). To obtain this flow, current approaches [3, 30] approximate $\phi_{0 \rightarrow t}$ as the following linear interpolation of the flow or latent vector:

$$\phi_{0 \rightarrow t}^{\theta^*} := t \cdot \phi_{0 \rightarrow 1}^{\theta^*} \quad \text{or} \quad \phi_{0 \rightarrow t}^{\theta^*} := \phi_{0 \rightarrow 1}^{t \cdot \theta^*}, \quad (5)$$

where $t \cdot \theta$ indicates the linear multiplication of latent vector. Therefore, the target I_t can be obtained by approximating it as $\hat{I}_t := I_0 \circ \phi_{0 \rightarrow t}^{\theta^*}$.

Limitations. As detailed in the latter part of Sec. 2.1, existing post-hoc linear interpolation approaches encounter two major challenges: firstly, they are prone to spatial distortion since the underlying network \mathcal{F}^{θ^*} that \hat{I}_t relies on does not account for the structural smoothness between two samples during network training; and secondly, they often suffer from artifacts resulting from the warping procedure. Moreover, the methods heavily rely on post-hoc linear multiplication, leading to potential overfitting to the linear assumption. In other words, these methods assume that the structures within a given 4D medical image move only in a linear direction, and the magnitude of this movement is linearly proportional to time.

4. Method

We introduce our Unsupervised Volumetric Interpolation Network, referred to as UVI-Net. The network first generates intermediate images and then employs cycle consistency constraints to reconstruct authentic images from these synthesized ones. In Sec. 4.1, we provide an overview and a detailed presentation of our method. The training and inference procedures are outlined in Sec. 4.2 and Sec. 4.3, respectively. Additionally, in Sec. 4.4, we introduce an instance-specific optimization method to further enhance our model's performance.

4.1. Methodology overview

To achieve a result exhibiting improved smoothness for the intermediate sample derived from the network, it is imperative for the network to access the pertinent information to

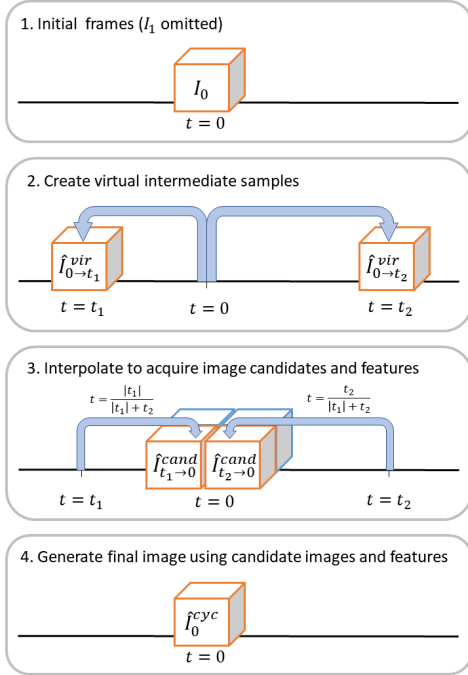


Figure 1. An overview of time-domain cycle consistency constraint. This image illustrates the process of generating \hat{I}_0^{cyc} . (1) I_0 and I_1 are given two input frames, with I_1 omitted for sake of readability. (2) We first generate virtual intermediate frames, and (3) subsequently generate back the frames with multi-resolution features (denoted as blue cubics). (4) The resulting reconstructed images \hat{I}_0^{cyc} must match the original input frame, I_0 .

the intermediate sample during the learning process. In light of this, we propose the cyclic structure model, which first generates the intermediate images and reconstructs them back to the two given input images. To ensure consistency and coherence in the generated images, we impose constraints of cycle consistency \mathcal{L}_{cyc} between the reconstructed samples, denoted as \hat{I}_0^{cyc} , \hat{I}_1^{cyc} , and the corresponding original samples I_0 , I_1 . The flow of the intermediate frame is then estimated using our flow calculation model with the parameter θ^* as follows:

$$\phi_{0 \rightarrow t}^{\theta^*} := t \cdot \phi_{0 \rightarrow 1}^{\theta^*} \quad \text{s.t.} \quad (6)$$

$$\theta^* := \arg \min_{\theta} \min_{\omega, \psi} \sum_{(I_0, I_1) \in \mathcal{D}} \mathcal{L}_{warp}^{\theta}(I_0, I_1) + \mathcal{L}_{cyc}^{(\theta, \omega, \psi)}(I_0, \hat{I}_0^{cyc}) + \mathcal{L}_{cyc}^{(\theta, \omega, \psi)}(I_1, \hat{I}_1^{cyc}), \quad (7)$$

where θ , ω , and ψ indicate the parameters for the flow calculation, feature extraction, and reconstruction models, which will be described in the below sections.

Unlike the current approach in Eq. (4), we allow the network to access intermediate samples and update them in its training, as described in Eq. (7), resulting in improved natural voxels. We first explain the process of obtaining \hat{I}_0^{cyc}

and provide a detailed explanation of \mathcal{L}_{cyc} in the following sections.

4.2. Training

The overall acquire procedure of \hat{I}_0^{cyc} and \hat{I}_1^{cyc} is illustrated in Fig. 1. First, we generate multiple virtual intermediate samples (see Step 2 in Fig. 1) by randomly sampling values of t_1 , t_2 , and t_3 as below.

$$\hat{I}_{t_1}^{vir} := I_0 \circ (t_1 \cdot \phi_{0 \rightarrow 1}^{\theta}) \quad -0.5 \leq t_1 \leq 0 \quad (8)$$

$$\hat{I}_{t_2}^{vir} := \begin{cases} I_0 \circ (t_2 \cdot \phi_{0 \rightarrow 1}^{\theta}) & 0 \leq t_2 \leq 0.5 \\ I_1 \circ ((1 - t_2) \cdot \phi_{1 \rightarrow 0}^{\theta}) & 0.5 \leq t_2 \leq 1 \end{cases} \quad (9)$$

$$\hat{I}_{t_3}^{vir} := I_1 \circ ((1 - t_3) \cdot \phi_{1 \rightarrow 0}^{\theta}) \quad 1 \leq t_3 \leq 1.5 \quad (10)$$

Since $\hat{I}_{t_1}^{vir}$ and $\hat{I}_{t_3}^{vir}$ are generated outside the time range between the two frames, we limit the maximum time offset to 0.5 to mitigate the occurrence of artifacts. When generating the $\hat{I}_{t_2}^{vir}$, a synthesized image between the two input images, we adopt the result created from the image—either I_0 or I_1 —that is closer to the reference point t_2 , to preserve the properties of the real image maximally.

Next, we interpolate the generated intermediate samples (see Step 3 in Fig. 1) to acquire the I_0 and I_1 's candidates as follows:

$$\hat{I}_{t_1 \rightarrow 0}^{cand} := \hat{I}_{t_1}^{vir} \circ \left(\frac{-t_1}{t_2 - t_1} \cdot \phi_{t_1 \rightarrow t_2}^{\theta} \right), \quad (11)$$

$$\hat{I}_{t_2 \rightarrow 0}^{cand} := \hat{I}_{t_2}^{vir} \circ \left(\frac{t_2}{t_2 - t_1} \cdot \phi_{t_2 \rightarrow t_1}^{\theta} \right), \quad (12)$$

$$\hat{I}_{t_2 \rightarrow 1}^{cand} := \hat{I}_{t_1}^{vir} \circ \left(\frac{1 - t_2}{t_3 - t_2} \cdot \phi_{t_2 \rightarrow t_3}^{\theta} \right), \quad (13)$$

$$\hat{I}_{t_3 \rightarrow 1}^{cand} := \hat{I}_{t_3}^{vir} \circ \left(\frac{t_3 - 1}{t_3 - t_2} \cdot \phi_{t_3 \rightarrow t_2}^{\theta} \right). \quad (14)$$

While warping the virtual frames, we simultaneously warp the feature space of the frames across multiple resolutions, obtaining a set of warped feature maps: $\mathcal{S}_{t_1 \rightarrow 0}$, $\mathcal{S}_{t_2 \rightarrow 0}$, $\mathcal{S}_{t_2 \rightarrow 1}$, and $\mathcal{S}_{t_3 \rightarrow 1}$. Specifically, following the architecture of our feature extractor as shown in Fig. 3, we extract feature maps resized to 1, 0.5, and 0.25 times their original size. Then, using the same optical flow as described in Eq. (11) to (14) (or downsampled as necessary), we obtain the final warped feature maps. This method enhances the reconstruction model's ability to make more accurate predictions by providing access to both voxel and feature information. Furthermore, we extract image representations at various levels, which have proven effective in previous research on video-related tasks [26, 29, 49].

Using these warped images and features, we obtain the predictions \hat{I}_0^{cyc} and \hat{I}_1^{cyc} using the reconstruction model \mathcal{R}^{ψ} (see Step 4 in Fig. 1). The model takes the distance-based weighted sum images and warped feature map sets,

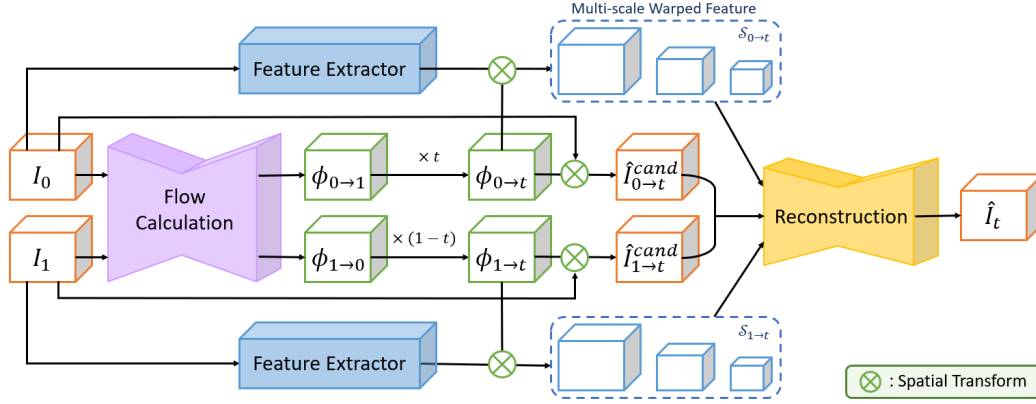


Figure 2. Schematic overview of our entire inference process. Starting with two input frames, I_0 and I_1 , we input the frames into the flow calculation model to obtain the approximated flow fields $\phi_{0 \rightarrow t}$ and $\phi_{1 \rightarrow t}$. We then warp the two frames using the obtained flow field, and similarly warp the multi-scale voxelwise features. Finally, we refine the distance-inversely weighted added image considering the information from multi-scale features, resulting in the final interpolated frame \hat{I}_t .

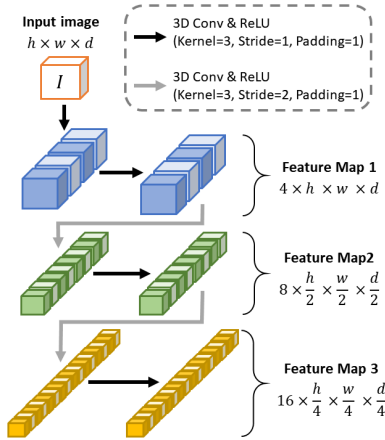


Figure 3. Architecture of the feature extractor module based on 3D Convolutional Neural Network (CNN). h , w , and d are the input image's height, width, and depth, respectively.

and reconstructs the original frames through residual corrections. Each element of the input feature map sets is fed into individual encoder layers of the reconstruction model and concatenated channel-wise. The procedure of the reconstruction model can be written as:

$$\hat{I}_0^{cyc} := \mathcal{R}^\psi(\hat{I}_{t_1 \rightarrow 0}^{cand} \oplus \hat{I}_{t_2 \rightarrow 0}^{cand}, \mathcal{S}_{t_1 \rightarrow 0}, \mathcal{S}_{t_2 \rightarrow 0}), \quad (15)$$

$$\hat{I}_1^{cyc} := \mathcal{R}^\psi(\hat{I}_{t_2 \rightarrow 1}^{cand} \oplus \hat{I}_{t_3 \rightarrow 1}^{cand}, \mathcal{S}_{t_2 \rightarrow 1}, \mathcal{S}_{t_3 \rightarrow 1}), \quad (16)$$

where \oplus indicates distance-based addition.

With the reconstructed images \hat{I}_0^{cyc} and \hat{I}_1^{cyc} , we can introduce the cycle consistency loss. Our cycle-consistent framework reconstructs real images from the generated intermediate images, thereby enhancing the smoothness of the interpolated images. Without time notation for clarity, consider the reconstructed image (\hat{I}^{cyc}) and corresponding real image (I). The cycle consistency loss is defined as:

$$\mathcal{L}_{cyc}^{(\theta, \omega, \psi)}(I, \hat{I}^{cyc}) = \mathcal{L}_{image}(I, \hat{I}^{cyc}) + \mathcal{L}_{reg}(\mathcal{R}^\psi), \quad (17)$$

where \mathcal{L}_{image} follows Eq. (3), and \mathcal{L}_{reg} acts as an L1 regularization term applied to the predicted residual of the reconstruction model. This term helps control excessive modification during the reconstruction process.

In essence, even without any intermediate frames, we utilize the given authentic frames as pseudo supervision for the intermediate frame, facilitated by the initially generated virtual intermediate samples \hat{I}^{vir} . By incorporating a cycle consistency constraint between the reconstructed and original authentic images, our approach enhances spatial continuity between the two images and generates high-quality virtual intermediate samples.

4.3. Inference

We illustrate the overall inference procedure of UVI-Net in Fig. 2. First, we obtain two optical flow $\phi_{0 \rightarrow 1}^{\theta^*}$ and $\phi_{1 \rightarrow 0}^{\theta^*}$, where θ^* follows Eq. (7). Next, we attain two I_t 's candidate as follows:

$$\hat{I}_{0 \rightarrow t}^{cand} := I_0 \circ \phi_{0 \rightarrow t} = I_0 \circ (t \cdot \phi_{0 \rightarrow 1}^{\theta^*}) \quad (18)$$

$$\hat{I}_{1 \rightarrow t}^{cand} := I_1 \circ \phi_{1 \rightarrow t} = I_1 \circ ((1-t) \cdot \phi_{1 \rightarrow 0}^{\theta^*}). \quad (19)$$

Finally, by reconstructing the final image with the two candidates considering the temporal distance, we derive \hat{I}_t as

$$\hat{I}_t := \mathcal{R}^\psi(\hat{I}_{0 \rightarrow t}^{cand} \oplus \hat{I}_{1 \rightarrow t}^{cand}, \mathcal{S}_{0 \rightarrow t}, \mathcal{S}_{1 \rightarrow t}), \quad (20)$$

where $\mathcal{S}_{0 \rightarrow t}$ and $\mathcal{S}_{1 \rightarrow t}$ are warped feature map sets from I_0 and I_1 , respectively. Remarkably, while baseline approaches can only use one of $\hat{I}_{0 \rightarrow t}$ and $\hat{I}_{1 \rightarrow t}$, we can engage both information and make to be symmetric even the order of I_0 and I_1 is switched.

4.4. Instance-Specific Optimization

Instance-specific optimization is a technique used to enhance the final performance by fine-tuning models for each

Dataset	Supervised	Method	PSNR \uparrow	NCC \uparrow	SSIM \uparrow	NMSE \downarrow	LPIPS \downarrow
Cardiac	✓	SVIN [16]	32.51 \pm 0.254	0.559 \pm 0.007	0.972 \pm 0.001	2.930 \pm 0.155	1.535 \pm 0.043
		MPVF [68]	33.15 \pm 0.238	0.561 \pm 0.006	0.971 \pm 0.001	2.435 \pm 0.133	1.941 \pm 0.055
		VM [3]	31.02 \pm 0.272	0.555 \pm 0.006	0.966 \pm 0.002	4.254 \pm 0.261	1.772 \pm 0.064
	✗	TM [10]	30.45 \pm 0.280	0.547 \pm 0.006	0.958 \pm 0.002	4.826 \pm 0.278	2.083 \pm 0.078
		Fourier-Net+ [24]	29.98 \pm 0.287	0.544 \pm 0.006	0.957 \pm 0.002	5.503 \pm 0.314	2.008 \pm 0.077
		R2Net [27]	28.59 \pm 0.278	0.509 \pm 0.007	0.930 \pm 0.003	7.281 \pm 0.329	3.482 \pm 0.138
		DDM [30]	29.71 \pm 0.221	0.541 \pm 0.006	0.956 \pm 0.002	5.007 \pm 0.239	2.136 \pm 0.066
		IDIR* [69]	31.56 \pm 0.275	0.557 \pm 0.006	0.968 \pm 0.001	3.806 \pm 0.249	1.675 \pm 0.061
		Ours (w/o inst opt.)	<u>33.57</u> \pm 0.275	0.565 \pm 0.007	<u>0.977</u> \pm 0.001	<u>2.409</u> \pm 0.159	<u>1.134</u> \pm 0.044
		Ours (w/ inst opt.)	33.59 \pm 0.268	0.565 \pm 0.007	0.978 \pm 0.001	2.384 \pm 0.157	1.066 \pm 0.041
Lung	✓	SVIN [16]	30.99 \pm 0.309	0.312 \pm 0.002	0.973 \pm 0.002	0.852 \pm 0.063	2.182 \pm 0.093
		MPVF [68]	31.18 \pm 0.344	0.310 \pm 0.003	0.972 \pm 0.002	0.761 \pm 0.075	2.554 \pm 0.092
		VM [3]	32.29 \pm 0.314	0.316 \pm 0.002	0.977 \pm 0.001	0.641 \pm 0.052	2.063 \pm 0.108
	✗	TM [10]	30.92 \pm 0.290	0.313 \pm 0.002	0.973 \pm 0.001	0.786 \pm 0.050	2.746 \pm 0.113
		Fourier-Net+ [24]	30.26 \pm 0.314	0.308 \pm 0.003	0.971 \pm 0.002	0.959 \pm 0.061	2.615 \pm 0.125
		R2Net [27]	29.34 \pm 0.270	0.294 \pm 0.003	0.962 \pm 0.002	1.061 \pm 0.051	3.277 \pm 0.122
		DDM [30]	30.37 \pm 0.271	0.308 \pm 0.003	0.971 \pm 0.002	0.905 \pm 0.065	2.283 \pm 0.106
		IDIR* [69]	32.91 \pm 0.309	0.321 \pm 0.003	0.980 \pm 0.002	0.586 \pm 0.055	2.035 \pm 0.112
		Ours (w/o inst opt.)	<u>33.90</u> \pm 0.382	0.319 \pm 0.003	0.980 \pm 0.002	<u>0.558</u> \pm 0.055	<u>1.512</u> \pm 0.112
		Ours (w/ inst opt.)	34.00 \pm 0.387	<u>0.320</u> \pm 0.003	0.980 \pm 0.002	0.552 \pm 0.055	1.489 \pm 0.093

Table 2. Quantitative comparison of interpolation results. These metrics were evaluated after repeating each experiment three times and collecting all frames. The model marked with an ‘*’ is trained exclusively on the test set, as it is designed for training on a single data pair only. For our model, the results with or without instance-specific optimization are both reported. The table presents both the average and standard deviation for each metric. NMSE and LPIPS values are presented in units of 10^{-2} . The best and second-best results for each metric are indicated with **bold** and underlined, respectively.

test sample. This approach was introduced by Balakrishnan et al. [2] within the unsupervised medical image warping domain. Despite our work being in a different task, this strategy remains applicable. Utilizing a model weight pre-trained on the training data, we fine-tune the model for a relatively small number of epochs on each test data. Such an adaptive approach is particularly beneficial in medical imaging, allowing for more personalized and accurate frame interpolation tailored to individual scans.

5. Experiments

This section describes the benchmark datasets for 4D medical imaging used in this study in Sec. 5.1. Next, Sec. 5.2 outlines some settings, including training details and metrics for performance evaluation. The results are comprehensively presented in Sec. 5.3, highlighting our method’s effectiveness and efficiency.

5.1. Datasets

To evaluate the performance of image interpolation, two 4D image datasets are used, each for the heart and lung. The ACDC cardiac dataset [6] consists of 100 4D temporal cardiac MRI images. End-diastolic and end-systolic phase images are used as the start and end images, respectively. The initial 90 alphabetically sorted samples form the training set, with the remaining used for the test set. The 4D-Lung dataset [22] consists of 82 chest CT scans for radiotherapy planning from 20 lung cancer patients. In each 4D-CT

study, the end-inspiratory (0% phase) and end-expiratory (50% phase) phase scans are set as the initial and final images, respectively. The first 68 CT scans from 18 patients in the dataset are included in the training set. For additional information for dataset, please refer to Appendix A.

5.2. Experimental settings

5.2.1 Baselines

For comparison with our proposed methods, six models are included as the baselines. VoxelMorph (VM) [2], TransMorph (TM) [10], Fourier-Net+ [24] and R2Net [27] are first initially trained with the provided dataset to calculate optical flow. Interpolated images are then obtained by linear scaling the optical flow, i.e., $t \cdot \phi_{0 \rightarrow t}^{\theta^*}$. Diffusion Deformable Model (DDM) [30] also uses dataset training but interpolates by scaling the latent vector, i.e., $\phi_{0 \rightarrow t}^{t \cdot \theta^*}$. For IDIR [69], it is crucial to clarify that it requires individual training for each target registration pair, leading to limited generalization, whereas our method is trained using a distinct training set and subsequently applied for inference on the target pairs. We also compared the results of our model with two supervised methods proposed for video interpolation on 4D medical images: SVIN [16] and MPVF [68]. Detailed information about the baseline models is in Appendix B.

5.2.2 Evaluation metrics

To evaluate the similarity between the predicted and ground truth images, metrics including PSNR (Peak Signal-to-

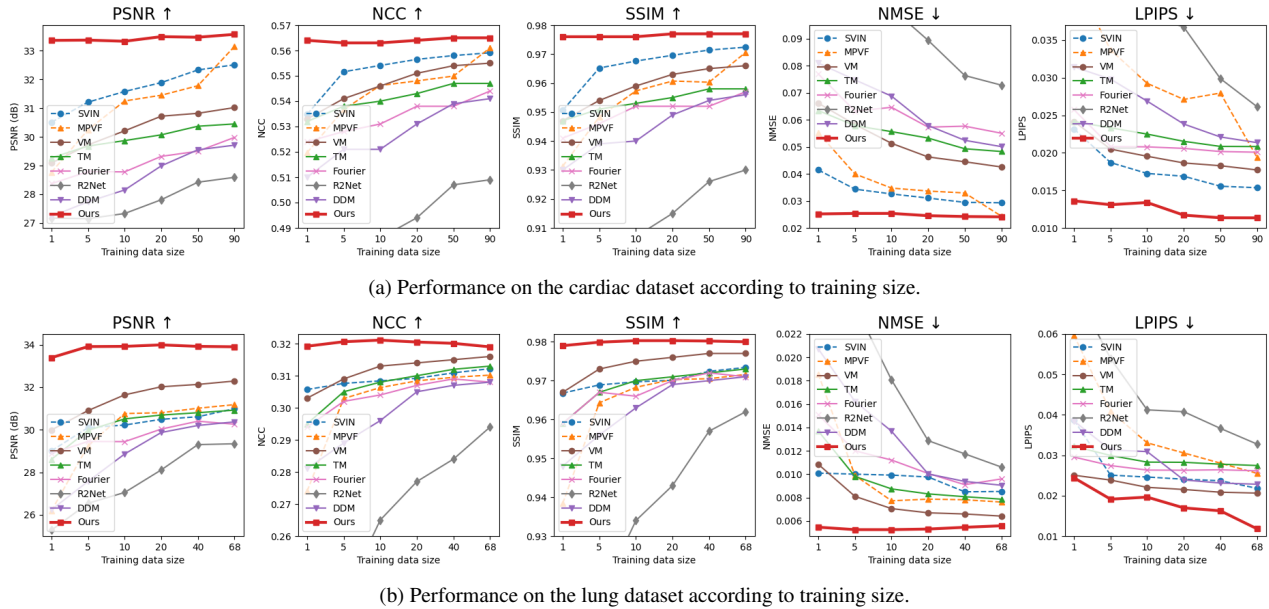


Figure 4. Performance trends based on the size of the training datasets. The dashed line represents a supervised setting. As depicted in this figure, we observe that the performance gap between our model and the baselines increases regardless of whether the setting is supervised or not, and irrespective of the dataset type. This demonstrates our model’s robustness, particularly in addressing data scarcity issues common in the medical domain.

Noise Ratio) [14], NCC (Normalized Cross Correlation), SSIM (Structural Similarity Index Measure) [79], NMSE (Normalized Mean Squared Error) and LPIPS (Learned Perceptual Image Patch Similarity) [76] are used. Since LPIPS is available only for 2D, it was averaged across slices along the x , y , and z axes. Each metric represents the voxel-wise similarity, correlation, structural similarity, reconstruction error, and perceptual similarity between the synthesized and authentic images.

5.2.3 Training details

For the flow calculation model, we employed the network designed in VoxelMorph [2]. As for the reconstruction model, we used a small size of 3D-UNet. The detailed configuration of the network and more details are described in Appendix C. The proposed method was implemented with PyTorch [52] using an NVIDIA Tesla V100 GPU. The training process takes approximately 4 hours for the cardiac dataset and 8 hours for the lung dataset, respectively. Instance-specific optimization took about 1.12 minutes per sample for ACDC and 3.13 minutes for 4D-Lung.

5.3. Results

5.3.1 Interpolation Result

The performance of interpolation compared to unsupervised and supervised methods is shown in Tab. 2. Our method consistently demonstrates superior performance among all

the models, outperforming others with a significant margin in every evaluation metric. This trend is observed across both heart and lung datasets, even in the absence of instance-specific optimization.

It is important to note that our approach surpasses IDIR [69], serving as a rigorous comparison baseline for our method due to IDIR’s test set-specific optimization. The core methodology behind IDIR undergoes unique adaptation for each test set pair, which involves retraining for every new instance. While this strategy enables IDIR to tailor its performance to each dataset, it restricts its practical applicability. Nevertheless, our method demonstrates substantial superiority over IDIR in terms of performance.

Supervised Models. Notably, our approach also surpasses supervised methods. An interesting observation is the varying performance of these supervised models across different datasets. As detailed in Tab. 1, the ACDC dataset contains significantly more frames compared to the 4D-Lung dataset. This discrepancy implies that the 4D-Lung dataset experiences limitations in terms of supervision quality. Therefore, the performance gap is more pronounced in the lung dataset, underscoring a critical insight: supervised models tend to underperform with limited supervision from intermediate frames. This pattern reaffirms the importance of our method’s ability to achieve high accuracy in scenarios with constrained supervision, highlighting its robustness and effectiveness in 4D medical VFI tasks.




Figure 5. Visualization examples from 4D cardiac and lung datasets. The model marked with an “*” is trained exclusively on the test set, while models marked with ‘(SL)’ are trained using supervised learning. Our method generates intermediate frames that are not only visually appealing but also precise, successfully retaining fine details and maintaining the structural integrity of the original images.

5.3.2 Effect of training dataset size

Fig. 4 illustrates the interpolation performance based on the number of training samples. With the test sets remaining fixed, the sizes of the training sets are reduced from their full down to one. Compared to the other five unsupervised baselines (VM, TM, Fourier-Net+, R2Net, and DDM), our method consistently exhibits superior performance across varying training set sizes, with performance gaps widening as the dataset size decreases. Remarkably, even with a minimal size comprising *only one sample*, our approach frequently outperforms the baselines that utilize the maximum training set size. It should be noted that IDIR is not included in this comparison, as it does not follow a traditional training process on a training set. For the two supervised baseline models (SVIN, MPVF), the performance also diminishes as the number of samples for supervision decreases, leading to an increasing performance gap between them and our model. Our consistent performance in scenarios with small datasets underscores the strengths of our approach in mitigating the challenges posed by data scarcity in the medical field.

5.3.3 Qualitative analysis

The comparison of qualitative results between interpolation methods is shown in Fig. 5. Our method consistently produces visually appealing and accurate intermediate frames, capturing fine details and preserving the structural integrity of the original images.

5.3.4 Downstream task

We also demonstrate that our interpolation method can be applied to downstream tasks. Specifically, we tested its effectiveness on segmentation data, which is relatively com-

plex, demonstrating our approach’s potential for augmenting 3D medical datasets. Details of the experimental setup and performance can be found in Appendix D.

5.3.5 Additional experiments

Additional experiments, including ablation studies and further qualitative results, are detailed in Appendix E. Moreover, we have analyzed the results of extrapolation to ensure that generated images during the training process do not exhibit any unnatural changes or issues.

6. Conclusion

Our framework, UVI-Net, effectively tackles the challenge of generating intermediate frames for 4D medical images through unsupervised volumetric interpolation. By leveraging pseudo supervision within a cyclic structure, our method ensures spatial continuity between the generated intermediate and real images. Experimental results on benchmark datasets validate the efficacy of our approach, revealing substantial improvements in intermediate frame quality across various evaluation metrics, surpassing both unsupervised and supervised baselines. Furthermore, our method has demonstrated robustness not only in situations of frame scarcity but also in data scarcity contexts. Ultimately, this study underscores the promise of unsupervised 3D flow-based interpolation and opens new avenues for research and development in the field of medical imaging.

Acknowledgement This study was supported by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIT) (No.2019-0-00075 Artificial Intelligence Graduate School Program (KAIST)) and Medical Scientist Training Program from the Ministry of Science & ICT of Korea.

References

- [1] Abolfazl Abdollahi, Biswajeet Pradhan, and Abdullah Alamri. Vnet: An end-to-end fully convolutional neural network for road extraction from high-resolution remote sensing data. *IEEE Access*, 8:179424–179436, 2020. [15](#)
- [2] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. An unsupervised learning model for deformable medical image registration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9252–9260, 2018. [2](#), [6](#), [7](#)
- [3] Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John Guttag, and Adrian V. Dalca. Voxelmorph: A learning framework for deformable medical image registration. *IEEE Transactions on Medical Imaging*, 38(8):1788–1800, 2019. [3](#), [6](#), [13](#), [14](#), [15](#)
- [4] M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61:139–157, 2005. [2](#)
- [5] J Bellec, F Arab-Ceschia, J Castelli, C Lafond, and E Chajon. Itv versus mid-ventilation for treatment planning in lung sbrrt: a comparison of target coverage and ptv adequacy by using in-treatment 4d cone beam ct. *Radiation Oncology*, 15:1–10, 2020. [1](#)
- [6] Olivier Bernard, Alain Lalande, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging*, 37(11):2514–2525, 2018. [2](#), [6](#)
- [7] Rhydian Caines, Naomi K Sisson, and Carl G Rowbottom. 4dct and vmat for lung patients with irregular breathing. *Journal of Applied Clinical Medical Physics*, 23(1):e13453, 2022. [1](#)
- [8] Xiaohuan Cao, Jianhuan Yang, Li Wang, Zhong Xue, Qian Wang, and Dinggang Shen. Deep learning based inter-modality image registration supervised by intra-modality similarity. In *Machine Learning in Medical Imaging: 9th International Workshop, MLMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 9*, pages 55–63. Springer, 2018. [2](#)
- [9] Pierre Charbonnier, Laure Blanc-Feraud, Gilles Aubert, and Michel Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st international conference on image processing*, pages 168–172. IEEE, 1994. [3](#)
- [10] Junyu Chen, Eric C Frey, Yufan He, William P Segars, Ye Li, and Yong Du. Transmorph: Transformer for unsupervised medical image registration. *Medical image analysis*, 82:102615, 2022. [6](#), [13](#), [15](#)
- [11] Zeyuan Chen, Yinbo Chen, Jingwen Liu, Xingqian Xu, Vidit Goel, Zhangyang Wang, Humphrey Shi, and Xiaolong Wang. Videoinr: Learning video implicit neural representation for continuous space-time super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2047–2057, 2022. [2](#)
- [12] Xianjin Dai, Yang Lei, Yabo Fu, Walter J Curran, Tian Liu, Hui Mao, and Xiaofeng Yang. Multimodal mri synthesis using unified generative adversarial networks. *Medical physics*, 47(12):6343–6354, 2020. [3](#)
- [13] Lee R Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945. [15](#)
- [14] R. Dosselmann and Xue Dong Yang. Existing and emerging image quality metrics. In *Canadian Conference on Electrical and Computer Engineering, 2005.*, pages 1906–1913, 2005. [7](#)
- [15] Dagmar Grob, Luuk Oostveen, Jan Rühaak, Stefan Heldmann, Brian Mohr, Koen Michielsen, Sabrina Dorn, Mathias Prokop, Marc Kachelrieß, Monique Brink, et al. Accuracy of registration algorithms in subtraction ct of the lungs: A digital phantom study. *Medical physics*, 46(5):2264–2274, 2019. [13](#)
- [16] Yuyu Guo, Lei Bi, Euijoon Ahn, Dagan Feng, Qian Wang, and Jinman Kim. A spatiotemporal volumetric interpolation network for 4d dynamic medical image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4726–4735, 2020. [2](#), [6](#), [13](#)
- [17] Yuyu Guo, Lei Bi, Dongming Wei, Liyun Chen, Zhengbin Zhu, Dagan Feng, Ruiyan Zhang, Qian Wang, and Jinman Kim. Unsupervised landmark detection-based spatiotemporal motion estimation for 4-d dynamic medical images. *IEEE Transactions on Cybernetics*, 2021. [2](#)
- [18] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 574–584, 2022. [15](#)
- [19] Weihua He, Kaichao You, Zhendong Qiao, Xu Jia, Ziyang Zhang, Wenhui Wang, Huchuan Lu, Yaoyuan Wang, and Jianxing Liao. Timereplayer: Unlocking the potential of event cameras for video interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17804–17813, 2022. [2](#)
- [20] Alessa Hering, Lasse Hansen, Tony CW Mok, Albert CS Chung, Hanna Siebert, Stephanie Häger, Annkristin Lange, Sven Kuckertz, Stefan Heldmann, Wei Shao, et al. Learn2reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. *IEEE Transactions on Medical Imaging*, 2022. [15](#)
- [21] Kan N Hor, Rolf Baumann, Gianni Pedrizzetti, Gianni Tonti, William M Gottliebson, Michael Taylor, D Woodrow Benson, and Wojciech Mazur. Magnetic resonance derived myocardial strain assessment using feature tracking. *JoVE (Journal of Visualized Experiments)*, (48):e2356, 2011. [1](#)
- [22] Geoffrey D. Hugo, Elisabeth Weiss, William C. Sleeman, Salman Balik, Paul J. Keall, Jun Lu, and Jeffrey F. Williamson. Data from 4d lung imaging of nscl patients. 2016. [2](#), [6](#)
- [23] Mi-Young Jeung, Philippe Germain, Pierre Croisille, Soraya El ghannudi, Catherine Roy, and Afshin Gangi. Myocardial tagging with mr imaging: overview of normal and pathologic findings. *Radiographics*, 32(5):1381–1398, 2012. [1](#)

- [24] Xi Jia, Alexander Thorley, Alberto Gomez, Wenqi Lu, Dipak Kotecha, and Jinming Duan. Fourier-net+: Leveraging band-limited representation for efficient 3d medical image registration. *arXiv preprint arXiv:2307.02997*, 2023. 2, 6, 13, 15
- [25] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slomo: High quality estimation of multiple intermediate frames for video interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9000–9008, 2018. 2, 3
- [26] Xin Jin, Longhai Wu, Jie Chen, Youxin Chen, Jayoon Koo, and Cheul-hee Hahm. A unified pyramid recurrent network for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1578–1587, 2023. 2, 3, 4
- [27] Ankita Joshi and Yi Hong. R2net: Efficient and flexible diffeomorphic image registration using lipschitz continuous residual networks. *Medical Image Analysis*, 89:102917, 2023. 2, 6, 13, 15
- [28] Neerav Karani, Lin Zhang, Christine Tanner, and Ender Konukoglu. An image interpolation approach for acquisition time reduction in navigator-based 4d mri. *Medical image analysis*, 54:20–29, 2019. 2
- [29] Rezaul Karim, He Zhao, Richard P Wildes, and Mennatullah Siam. Med-vt: Multiscale encoder-decoder video transformer with application to object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6323–6333, 2023. 4
- [30] Boah Kim and Jong Chul Ye. Diffusion deformable model for 4d temporal medical image generation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part I*, pages 539–548. Springer, 2022. 2, 3, 6, 13, 15
- [31] Boah Kim, Dong Hwan Kim, Seong Ho Park, Jieun Kim, June-Goo Lee, and Jong Chul Ye. Cyclemorph: cycle consistent unsupervised deformable image registration. *Medical image analysis*, 71:102036, 2021. 2
- [32] Boah Kim, Inhwa Han, and Jong Chul Ye. Diffusemorph: Unsupervised deformable image registration using diffusion model. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 347–364. Springer, 2022. 2
- [33] Taewoo Kim, Yujeong Chae, Hyun-Kurl Jang, and Kuk-Jin Yoon. Event-based video frame interpolation with cross-modal asymmetric bidirectional motion fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18032–18042, 2023. 3
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 14, 15
- [35] Dongyang Kuang. Cycle-consistent training for reducing negative jacobian determinant in deep registration networks. In *Simulation and Synthesis in Medical Imaging: 4th International Workshop, SASHIMI 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13, 2019, Proceedings 4*, pages 120–129. Springer, 2019. 2
- [36] seungmin Lee, Seongwook Yoon, and Sanghoon Sull. Unsupervised video frame interpolation using online refinement. In *Institute of Electronics, Information and Communication Engineers*, 2020. 1, 2
- [37] Sungho Lee, Narae Choi, and Woong Il Choi. Enhanced correlation matching based video frame interpolation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2839–2847, 2022. 2
- [38] Yang Lei, Yabo Fu, Tonghe Wang, Yingzi Liu, Pretesh Patel, Walter J Curran, Tian Liu, and Xiaofeng Yang. 4d-ct deformable image registration using multiscale unsupervised deep learning. *Physics in Medicine & Biology*, 65(8):085003, 2020. 2
- [39] Yu-Lun Liu, Yi-Tung Liao, Yen-Yu Lin, and Yung-Yu Chuang. Deep video frame interpolation using cyclic frame generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8794–8802, 2019. 1, 2
- [40] Yao Lu, Jack Valmadre, Heng Wang, Juho Kannala, Mehrtash Harandi, and Philip Torr. Devon: Deformable volume network for learning optical flow. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2705–2713, 2020. 2
- [41] Elliot McVeigh and Cengizhan Ozturk. Imaging myocardial strain. *IEEE Signal Processing Magazine*, 18(6):44–56, 2001. 1
- [42] Elliot R McVeigh, Amir Pourmorteza, Michael Guttman, Veit Sandfort, Francisco Contijoch, Suhas Budhiraja, Zhen-nong Chen, David A Bluemke, and Marcus Y Chen. Regional myocardial strain measurements from 4dct in patients with normal lv function. *Journal of cardiovascular computed tomography*, 12(5):372–378, 2018. 1
- [43] Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, and Christopher Schroers. Phasenet for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 498–507, 2018. 2
- [44] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016. 15
- [45] Kotaro Mizuno and Masahiro Muto. Preoperative evaluation of pleural adhesion in patients with lung tumors using four-dimensional computed tomography performed during natural breathing. *Medicine*, 100(47), 2021. 1
- [46] Dong Nie, Roger Trullo, Jun Lian, Caroline Petitjean, Su Ruan, Qian Wang, and Dinggang Shen. Medical image synthesis with context-aware generative adversarial networks. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part III 20*, pages 417–425. Springer, 2017. 3
- [47] Dong Nie, Roger Trullo, Jun Lian, Li Wang, Caroline Petitjean, Su Ruan, Qian Wang, and Dinggang Shen. Medical image synthesis with deep convolutional adversarial networks. *IEEE Transactions on Biomedical Engineering*, 65(12):2720–2730, 2018. 3

- [48] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1710, 2018. [2](#)
- [49] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5437–5446, 2020. [2](#), [4](#)
- [50] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *Proceedings of the IEEE international conference on computer vision*, pages 261–270, 2017. [2](#)
- [51] Junheum Park, Jintae Kim, and Chang-Su Kim. Biformer: Learning bilateral motion estimation via bilateral transformer for 4k video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1568–1577, 2023. [2](#), [3](#)
- [52] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. [7](#)
- [53] Tomer Peleg, Pablo Szekely, Doron Sabo, and Omry Sendik. Im-net for high resolution video frame interpolation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2398–2407, 2019. [2](#)
- [54] Fitsum A Reda, Deqing Sun, Aysegul Dundar, Mohammad Shoeybi, Guilin Liu, Kevin J Shih, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Unsupervised video interpolation using cycle consistency. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 892–900, 2019. [1](#), [2](#), [3](#)
- [55] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. [14](#), [15](#)
- [56] Elisabeth Sartoretti, Thomas Sartoretti, Christoph Binkert, Arash Najafi, Árpád Schwenk, Martin Hinnen, Luuk van Smoorenburg, Barbara Eichenberger, and Sabine Sartoretti-Schefer. Reduction of procedure times in routine clinical practice with compressed sense magnetic resonance imaging technique. *PLoS One*, 14(4):e0214887, 2019. [1](#)
- [57] Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. [15](#)
- [58] Yucheng Shu, Hao Wang, Bin Xiao, Xiuli Bi, and Weisheng Li. Medical image registration based on uncoupled learning and accumulative enhancement. In *Medical Image Computing and Computer Assisted Intervention 2021*, pages 3–13, Cham, 2021. Springer International Publishing. [3](#)
- [59] Hyeonjun Sim, Jihyong Oh, and Munchurl Kim. Xvfi: extreme video frame interpolation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14489–14498, 2021. [2](#)
- [60] Amber L Simpson, Michela Antonelli, Spyridon Bakas, Michel Bilello, Keyvan Farahani, Bram Van Ginneken, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, et al. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *arXiv preprint arXiv:1902.09063*, 2019. [15](#)
- [61] Li Siyao, Shiyu Zhao, Weijiang Yu, Wenxiu Sun, Dimitris Metaxas, Chen Change Loy, and Ziwei Liu. Deep animation video interpolation in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6587–6595, 2021. [2](#)
- [62] Pierre Soille. *Erosion and Dilation*, pages 63–103. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. [13](#)
- [63] Hessam Sokooti, Bob De Vos, Floris Berendsen, Boudewijn PF Lelieveldt, Ivana Išgum, and Marius Staring. Nonrigid image registration using multi-scale 3d convolutional neural networks. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11–13, 2017, Proceedings, Part I 20*, pages 232–239. Springer, 2017. [2](#)
- [64] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [2](#)
- [65] Shuo Chen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1279–1288, 2017. [2](#)
- [66] Lu Wang, Shelly Hayes, Kamen Paskalev, Lihui Jin, Mark K Buyyounouski, Charlie C-M Ma, and Steve Feigenberg. Dosimetric comparison of stereotactic body radiotherapy using 4d ct and multiphase ct images for treatment planning of lung cancer: evaluation of the impact on daily dose coverage. *Radiotherapy and Oncology*, 91(3):314–324, 2009. [1](#)
- [67] Wei-Hao Wang, Chia-Yu Sung, Shih-Chung Wang, and Yu-Hsuan Joni Shao. Risks of leukemia, intracranial tumours and lymphomas in childhood and early adulthood after pediatric radiation exposure from computed tomography. *CMAJ*, 195(16):E575–E583, 2023. [1](#)
- [68] Tzu-Ti Wei, Chin Kuo, Yu-Chee Tseng, and Jen-Jee Chen. Mpvf: 4d medical image inpainting by multi-pyramid voxel flows. *IEEE Journal of Biomedical and Health Informatics*, 2023. [3](#), [6](#), [13](#)
- [69] Jelmer M Wolterink, Jesse C Zwienenberg, and Christoph Brune. Implicit neural representations for deformable image registration. In *International Conference on Medical Imaging with Deep Learning*, pages 1349–1359. PMLR, 2022. [2](#), [6](#), [7](#), [13](#)
- [70] Mian Xi, Meng-Zhong Liu, Xiao-Wu Deng, Li Zhang, Xiao-Yan Huang, Hui Liu, Qiao-Qiao Li, Yong-Hong Hu, Ling Cai, and Nian-Ji Cui. Defining internal target volume (itv) for hepatocellular carcinoma using four-dimensional ct. *Radiotherapy and Oncology*, 84(3):272–278, 2007. [1](#)
- [71] Xiaoyu Xiang, Yapeng Tian, Yulun Zhang, Yun Fu, Jan P Allebach, and Chenliang Xu. Zooming slow-mo: Fast and

- accurate one-stage space-time video super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3370–3379, 2020. [2](#), [3](#)
- [72] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127:1106–1125, 2019. [2](#)
- [73] Heran Yang, Jian Sun, Aaron Carass, Can Zhao, Junghoon Lee, Zongben Xu, and Jerry Prince. Unpaired brain mr-to-ct synthesis using a structure-constrained cycleGAN. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 174–182. Springer, 2018. [3](#)
- [74] Xiao Yang, Roland Kwitt, Martin Styner, and Marc Niethammer. Quicksilver: Fast predictive image registration—a deep learning approach. *NeuroImage*, 158:378–396, 2017. [2](#)
- [75] Guozhen Zhang, Yuhan Zhu, Haonan Wang, Youxin Chen, Gangshan Wu, and Limin Wang. Extracting motion and appearance via inter-frame attention for efficient video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5682–5692, 2023. [2](#)
- [76] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [7](#)
- [77] Shengyu Zhao, Yue Dong, Eric I-Chao Chang, and Yan Xu. Recursive cascaded networks for unsupervised medical image registration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. [3](#)
- [78] Kun Zhou, Wenbo Li, Xiaoguang Han, and Jiangbo Lu. Exploring motion ambiguity and alignment for high-quality video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22169–22179, 2023. [2](#)
- [79] Wang Zhou. Image quality assessment: from error measurement to structural similarity. *IEEE transactions on image processing*, 13:600–613, 2004. [7](#)