

# Label-Efficient Data Augmentation with Video Diffusion Models for Guidewire Segmentation in Cardiac Fluoroscopy

Shaoyan Pan<sup>1\*</sup>, Yikang Liu<sup>2</sup>, Lin Zhao<sup>2</sup>, Eric Z. Chen<sup>2</sup>, Xiao Chen<sup>2</sup>, Terrence Chen<sup>2</sup>, Shanhui Sun<sup>2</sup>

<sup>1</sup>Emory University, Atlanta, GA, USA

<sup>2</sup>United Imaging Intelligence, Boston, MA, USA

pshaoya@emory.edu, {yikang.liu, lin.zhao01, zhang.chen, xiao.chen01, terrence.chen, shanhui.sun}@uii-ai.com

## Abstract

The accurate segmentation of guidewires in interventional cardiac fluoroscopy videos is crucial for computer-aided navigation tasks. Although deep learning methods have demonstrated high accuracy and robustness in wire segmentation, they require substantial annotated datasets for generalizability, underscoring the need for extensive labeled data to enhance model performance. To address this challenge, we propose the Segmentation-guided Frame-consistency Video Diffusion Model (SF-VD) to generate large collections of labeled fluoroscopy videos, augmenting the training data for wire segmentation networks. SF-VD leverages videos with limited annotations by independently modeling scene distribution and motion distribution. It first samples the scene distribution by generating 2D fluoroscopy images with wires positioned according to a specified input mask, and then samples the motion distribution by progressively generating subsequent frames, ensuring frame-to-frame coherence through a frame-consistency strategy. A segmentation-guided mechanism further refines the process by adjusting wire contrast, ensuring a diverse range of visibility in the synthesized image. Evaluation on a fluoroscopy dataset confirms the superior quality of the generated videos and shows significant improvements in guidewire segmentation.

**Extended version** — <https://arxiv.org/pdf/2412.16050>

## Introduction

Guidewires are frequently used in interventional surgeries to navigate blood vessels and guide other medical devices, such as catheters, stents, or balloons, to precise locations within the body. Guidewire segmentation in X-ray fluoroscopy videos is a critical task in computer-aided cardiac interventional procedures. It often serves as a prerequisite for functions such as 4D guidewire reconstruction (Vöth et al. 2023) and the co-registration of angiography and intravascular ultrasound images (Prasad et al. 2016). This task is inherently challenging due to factors such as image noise, complex scenes, projective foreshortening, object occlusion, rapid guidewire motion, and the thin, elongated shape of the guidewire.

\*The work conducted by the first author was carried out during an internship at United Imaging Intelligence.  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Accurate and sufficient annotated data are essential for training guidewire segmentation models, but labeling thin and elongated guidewires is both error-prone and labor-intensive. To address this challenge, this paper proposes a method to generate synthetic videos paired with guidewire masks as augmented data, aimed at enhancing the performance of downstream segmentation models (Fig. 1). We trained a generative diffusion model (DM) on a small labeled dataset along with a larger unannotated dataset, enabling the generation of multiple annotated fluoroscopy videos with varying backgrounds and wire appearances from a given semantic wire mask.

In medical data synthesis, Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) encounter challenges such as unstable training and limited video fidelity and diversity, particularly due to the complex distribution of medical images and small dataset sizes. Denoising Diffusion Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020; Song et al. 2020) have shown promise in generating diverse medical images (Yu et al. 2023; Du et al. 2023; Pan et al. 2023b). Similarly, video DDPMs (Blattmann et al. 2023b,a; Ho et al. 2022) demonstrate potential in synthesizing medical videos conditioned on semantic contents (Lugmayr et al. 2022; Wang et al. 2022; Du et al. 2023).

However, our task of synthesizing fluoroscopy videos for data augmentation in wire segmentation poses unique challenges and specific requirements. Firstly, it is challenging to acquire fluoroscopy images/videos at the scale of natural ones, and annotating wires in numerous videos is almost impractical. Therefore, it is beneficial to minimize the labeling effort while preserving the fidelity and variability of the synthesized videos. Secondly, since guidewires are extremely thin (typically 2-3 pixels wide) and tend to be occluded by other objects or noise, detection often relies on the contrast between fast-moving guidewires and the slower-moving background, which needs to be reflected in synthesized videos. Thirdly, the nuanced and varied appearance of wires against complex backgrounds in fluoroscopy videos requires a conditioning approach that accurately captures wire masks and generates diverse wire appearances. Inadequate representation of these masks compromises the fidelity and variability of wire depiction in the synthesized videos, reducing their value for segmentation tasks.

We propose a label-efficient data augmentation method




Figure 1: SF-VD as a strong data augmentation tool for training guidewire segmentation networks.

for guidewire segmentation in cardiac fluoroscopy videos, named Segmentation-guided Frame-consistency Video Diffusion Model (SF-VD). First, we use two separate 2D diffusion models (DMs) to independently learn scene distribution and cross-frame motion: the first model synthesizes a static image conditioned on a wire mask, while the second model generates subsequent frames based on the static image and the wire masks for those frames. This strategy allows the first model to effectively leverage data with limited labeled frames, while the second model ensures cross-frame consistency by training on fully annotated videos. Together, these models generate more diverse backgrounds than would be possible using only fully annotated video datasets. Additionally, the second model employs a “frame-consistency” sampling strategy that produces more realistic and diverse videos compared to 3D models, especially given our limited data. The second key idea is to modulate wire contrast and enhance the variability of wire appearance and visibility in synthetic frames. This is achieved by using a segmentation model to guide the diffusion reverse process, similar to the approach used in classifier guidance (Dhariwal and Nichol 2021).

We benchmarked the proposed method against various data augmentation methods including other video diffusion models, using several segmentation model architectures, which included both generic state-of-the-art (SOTA) models and those specifically designed for guidewire segmentation. Our results demonstrate that our method outperforms other data augmentation approaches and enhances the segmentation performance of all tested models.

In summary, the key contributions of this work are:

- Introduced SF-VD, a label-efficient video synthesis model that enhances downstream wire segmentation performance. To the best of our knowledge, this is the first work on fluoroscopy video synthesis and the first to use a generative model for data augmentation in guidewire segmentation.

- Proposed a two-model strategy to independently learn scene and motion distributions, addressing the challenges of limited medical data and guidewire annotations.
- Proposed a frame-consistency strategy to ensure coherent motion and content across fluoroscopy frames.
- Introduced diffusion guidance by a segmentation network to adjust guidewire contrast, improving visualization and diversity in fluoroscopy videos.

## Related Work

**Video Diffusion Model** Recent advancements in DMs have driven progress in video generation. Several studies (Ho et al. 2022; Ho et al. 2022; Liu et al. 2024) train 3D DMs to model video sequences as volumetric images. However, optimizing these models is challenging due to the high-dimensional nature of the data, which requires a large amount of video data and compute (Chen et al. 2024). In response, some researchers used pretrained 2D image DMs for consistent frame generation. They either integrated additional structures to fine-tune the pretrained image DMs (Yang, Srivastava, and Mandt 2023; Ge et al. 2023; Gu et al. 2023; Guo et al. 2023; Blattmann et al. 2023a), or employed a completely training-free approach (Khachatriyan et al. 2023; Wu et al. 2022), using the image DMs as a generative prior to generate videos. These methods reduce optimization challenges and improve video fidelity. However, directly applying these techniques to our problem is not optimal, since fine-tuning pretrained image DMs still requires a substantial number of labeled videos, which are costly to obtain for medical data. Moreover, precise pixel-level conditioning is crucial for our task, and maintaining this accuracy during sampling from a pretrained image DM is not straightforward. To address these challenges, we propose training a DM from scratch with pixel-wise wire conditioning and a two-model strategy, which independently models scene and motion distributions to reduce the amount of required data.

**Conditional Diffusion Model** Recent advancements in conditional DMs have significantly enhanced personalized, task-specific image and video generation by optimizing condition integration for better control. Key approaches include ‘network injection’ DMs, which embed conditions within the denoising network for controlled generation (Wang et al. 2022; Avrahami et al. 2023), and ‘process injection’ DDPMs, which modify the diffusion trajectory by incorporating conditions directly into the diffusion steps (Lugmayr et al. 2022; Mao, Wang, and Aizawa 2023; Liang et al. 2024). Classifier-guided-based DMs utilize an auxiliary network to preprocess conditions before injection, offering precise guidance (Dhariwal and Nichol 2021; Gafni et al. 2022; Zhang, Rao, and Agrawala 2023). Classifier-free DMs (Ho 2022; Nichol et al. 2021) employ a simpler architecture without explicit condition encoding, making them suitable for applications such as text-guided, segmentation-guided generation, and colorization. In the medical field, these methods have been adapted for conditional image generation, such as SegGuidedDiff (Konz et al. 2024) for generating breast MRI and abdominal CT images with organ segmentation maps, Med-DDPM (Dorjsembe et al. 2023) for

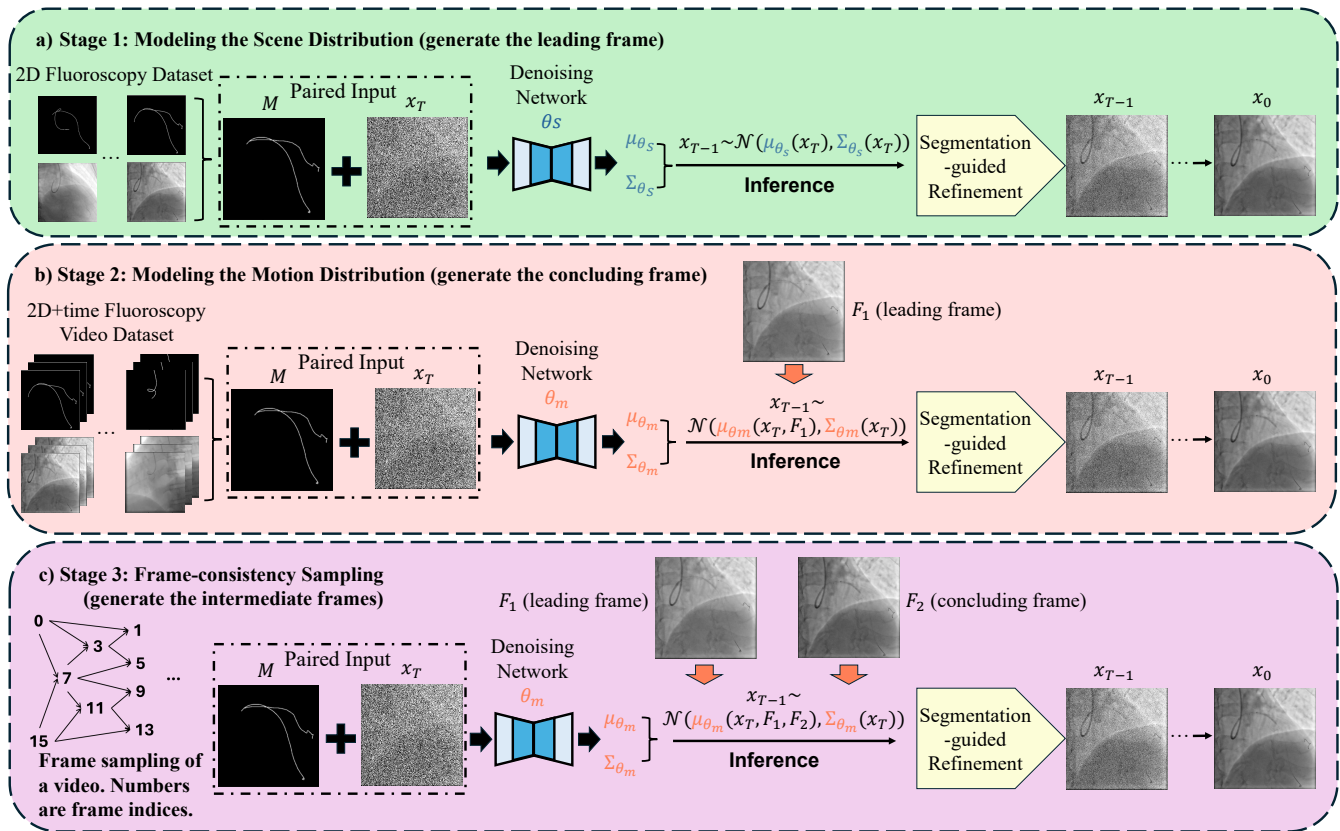


Figure 2: Overview of the three-stage approach in SF-VD.

generating 3D semantic brain MRI with brain organ segmentation maps, and ArSDM (Du et al. 2023) for producing colonoscopy images with polyps using polyp segmentation maps. These advancements reveal the potential of conditional DMs to generate medical images with precise, pixel-wise correspondence to specified segmentation maps. Building on this, our work introduces SF-VD, a model that uses wire masks to generate videos with clear wires precisely aligned to the given masks. In addition, SF-VD incorporates a controllable conditioning strategy that modulates wire contrast, enhancing the variability of wire appearance and visibility across synthetic frames.

**Guidewire Segmentation in Fluoroscopy Videos** Advancements in deep learning, particularly with convolutional neural networks (CNNs) (Milletari, Navab, and Ahmadi 2016; Isensee et al. 2020) and vision Transformers (Tang et al. 2022; Pan et al. 2023a), have shown promise in accurately segmenting medical objects, including organs and interventional devices (Liu et al. 2020; González, Bravo-Sánchez, and Arbelaez 2020). Innovative solutions such as pyramid attention recurrent networks (PARN) (Zhou et al. 2020), temporal Transformer network (TTN) (Zhang et al. 2021), and 4D recurrent networks (4D-RCN) (Lee et al. 2020) have been proposed to enhance feature capture across video frames, leading to improved guidewire segmentation. Additional strategies like real-time guidewire

network (RTGN) using a preliminary guidewire detection network for bounding box setup (Zhang et al. 2022), and Ariadne+ (Caporali et al. 2022), which applies graph theory for post-process refinement and false negative reduction, have been proposed. These models focus on development of neural network architectures and segmentation algorithms, but the scarcity of annotated training videos is overlooked, which often limits their performance on out-of-distribution videos. We propose a label-efficient data augmentation method using conditional DMs to improve segmentation network training with synthetic data that cover a wide range of guidewire contrast, background appearances, and motions.

## Methodology

**Overview** The aim of the proposed video generation method, SF-VD, is to take a sequence of guidewire masks (as illustrated in Fig. 1) and synthesize fluoroscopy videos with guidewires at the specified pixels, which exhibit good variation, fidelity, and cross-frame consistency. The method aims to efficiently utilize fluoroscopy videos and mask annotations, given their limited availability, and to serve as a data augmentation tool for training wire segmentation models.

To address of the problem of overfitting when directly modeling the complex condition distribution of fluoroscopy videos, we propose to decompose the distribution  $P(\mathbf{V}|\mathbf{M})$

$\mathbf{V} = \{I^{(0)} \dots I^{(N-1)}\}$  is a  $N$ -frame video and  $\mathbf{M} = \{M^{(0)} \dots M^{(N-1)}\}$  is its corresponding guidewire masks into two components: a scene distribution  $P(I^{(0)} | M^{(0)})$ , which accounts for factors such as imaging angles, imaging parameters, objects in view, and patient anatomy; and a motion distribution  $P(I^{(1)} \dots I^{(N-1)} | M^{(1)} \dots M^{(N-1)}, I^{(0)})$ , which captures object motion due to heartbeats and breathing. The strategy also helps to leverage the video with partially labeled frames, which can be used to learn the scene distribution.

We trained two 2D DMs to independently model each component and employed a three-stage approach (one stage to sample the scene distribution; two stages to sample the motion distribution) to synthesize videos, as shown in Fig. 2. In addition, a segmentation-guided mechanism was introduced to refine the diffusion reverse process throughout, ensuring controllable wire depiction across frames.

**Modeling the Scene Distribution** Modeling a distribution with DMs involve two primary Gaussian processes: forward and reverse diffusion (Nichol and Dhariwal 2021). In the forward process, an uncorrupted image  $x_0$  undergoes  $T$  diffusion steps, progressively incorporating noise  $\epsilon$ , resulting in a sequence of increasingly noisy frames:  $\{x_0, \dots, x_T\}$ . This process is modeled as a Gaussian process with a specified noise schedule  $\{\alpha_0, \dots, \alpha_T\}$ :  $x_t \sim \mathcal{N}(\sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\mathbf{I})$ , where  $\bar{\alpha}_t$  is the product of the noise coefficients up to time  $t$ ,  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ .

Next, a network parameterized by  $\theta$  is trained to reverse the forward process, starting from a noise sample and sequentially removing noise  $\epsilon$  to recover the original data.

To model the scene distribution of fluoroscopy videos, we want to generate an image conditioned on a guidewire mask. Given a noisy sample  $x_t$  and the time step  $t$ , the model predicts the previous less noisy sample  $x_{t-1}$  with the mask  $M$ :

$$x_{t-1} \sim p(x_t | M) = \mathcal{N}(\mu_\theta(x_t, t, M), \sigma_\theta^2(x_t, t, M)\mathbf{I})$$

$$\text{where } \mu_\theta(x_t, t, M) = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t, M) \right)$$

$$\text{and } \sigma_\theta^2(x_t, t, M) = \exp(v_\theta v_f + (1 - v_\theta)v_r) \quad (1)$$

Here, the condition mask  $M$  is channel-wise concatenated with  $x_t$ , and fed to the network along with the time step  $t$  to estimate the noise  $\epsilon_\theta$  and variance  $v_\theta$ .  $v_f$  and  $v_r$  are two variance-related constants defined in IDDPM (Nichol and Dhariwal 2021). The optimization for  $\mu_\theta$  and  $\sigma_\theta$  follows IDDPM, as shown in Appendix A.

The network was trained on a mixture of annotated and unannotated images to increase background variability. The strategy aligns with the formulation of the classifier-free (CF-) DDPM (Ho 2022), as the network simultaneously models both the unconditional and conditional image distributions. CF-DDPM approximates the conditional score function as:

$$\begin{aligned} \nabla_{x_t} \log p(x_t | M) &= \nabla_{x_t} \log p(x_t) + \nabla_{x_t} \log p(M | x_t) \\ &\approx \nabla_{x_t} \log p(x_t) + \omega(\nabla_{x_t} \log p(x_t | M) - \nabla_{x_t} \log p(x_t)) \end{aligned} \quad (2)$$

Unlike the original CF-DDPM where variance is constant, we incorporate a predicted variance to express the score function as:

$$\nabla_{x_t} \log p(x_t | M) = \frac{\epsilon_\theta(x_t, t, M)}{\sigma_\theta(x_t, t, M)} \quad (3)$$

Rewriting Eqs. 2 and 3 into the format of Eq. 1, we have the updated noise as:

$$\bar{\epsilon}_\theta(x_t, t, M) = (1 - \omega)\epsilon_\theta(x_t, t) + \omega\epsilon_\theta(x_t, t, M) \quad (4)$$

, where  $\omega$  is empirically selected as 0.7. For the variance  $\sigma_\theta^2(x_t, t)$ , we retain the original form  $\sigma_\theta^2(x_t, t, M)$  since the updated form does not improve performance, likely because the unconditional variance is difficult to estimate accurately.

**Modeling the Motion Distribution** To model the motion distribution, we propose to generate another frame based on a previously generated frame and a mask with a 2D DM. Note that modeling motion in cardiac fluoroscopy videos is typically simpler than in natural videos because, in fluoroscopy, objects tend to stay within the field of view and their appearance remains relatively stable. Therefore, a 2D conditional DM is sufficient for the task and helps prevent overfitting with limited training data.

Similar to Eq. 2, with the conditions of an initial frame  $F$  and a mask  $M$ , we have:

$$\begin{aligned} \nabla_{x_t} \log p(x_t | F, M) &= \nabla_{x_t} \log p(x_t | M) + \nabla_{x_t} \log p(F | x_t, M) \\ &\approx \nabla_{x_t} \log p(x_t | M) + \omega(\nabla_{x_t} \log p(x_t | F, M) \\ &\quad - \nabla_{x_t} \log p(x_t | M)) \end{aligned} \quad (5)$$

Similarly, the predicted variance in Eq. 1 is  $\sigma_\theta^2(x_t, t, F, M)$ , and the predicted noise is:

$$\bar{\epsilon}_\theta(x_t, t, M, F) = (1 - \omega)\epsilon_\theta(x_t, t, M) + \omega\epsilon_\theta(x_t, t, M, F) \quad (6)$$

**Frame-consistency Sampling** The networks described above enable us to generate the first frame of a video, followed by each subsequent frame sequentially. However, this approach produced videos with jittery background motion, negatively impacting segmentation performance (see Appendix D). To address this, we introduce ‘‘frame-consistency’’ sampling, where a frame is generated with the conditions of the preceding and following frames. Denoting these frames as  $F_1$  and  $F_2$ , similar to Eqs. 2 and 5, the conditional score function is:

$$\begin{aligned} \nabla_{x_t} \log p(x_t | F_1, F_2, M) &= \nabla_{x_t} \log p(x_t, F_1, F_2, M) - \nabla_{x_t} \log p(F_1, F_2, M) \\ &= \nabla_{x_t} \log p(x_t | M) + \nabla_{x_t} \log p(F_1 | x_t, M) \\ &\quad + \nabla_{x_t} \log p(F_2 | x_t, F_1, M) \end{aligned} \quad (7)$$

Similarly, the predicted noise is

$$\begin{aligned} \bar{\epsilon}_\theta(x_t, t, M, F_1, F_2) &= (1 - 2\omega)\epsilon_\theta(x_t, t, M) \\ &\quad + \omega\epsilon_\theta(x_t, t, M, F_1) + \omega\epsilon_\theta(x_t, t, M, F_2) \end{aligned} \quad (8)$$




Figure 3: Video Quality Comparison among SF-VD and Benchmarked Video Diffusion Models. Each sample includes a sequence of 4 frames, displayed from left to right, with the first frame on the left and the last frame on the right. The blue arrows indicate obvious diaphragm motions, and yellow arrows indicate less apparent motions




Figure 4: The Segmentation-guided Mechanism Generates Videos with Wires with Various Contrasts. Each row shows a video sequence from the first frame on the left to the last frame on the right.

**A Three-stage Approach** Taken together, as shown in Fig. 2, we first use the scene-distribution model  $\theta_s$  to generate the first (leading) frame, followed by the motion-distribution model  $\theta_m$  to generate the last (concluding) frame based on the leading frame. Next, we use the frame-

consistency sampling to generate the middle frame, conditioned on the leading and concluding frames. The remaining frames are generated in the same manner (the sampling order of a 16-frame video is shown on the left in Fig. 2c), by iteratively generating frames equidistant from the previously

generated ones until all frames are completed.

In practice, the motion-distribution model  $\theta_m$  also takes the frame distance between the generated frame and the conditional frame as inputs.  $\omega$  in Eqs. 6 and 8 were set to  $-2.5$  and  $-1.5$  when generating the concluding frame and intermediate frames, respectively.

**Segmentation-guided Mechanism** To modulate the wire contrast in fluoroscopy videos, a segmentation network  $\psi$ , trained on the 2D fluoroscopy dataset, was used as a sampling guidance. The segmentation-guided mechanism is integrated into the estimated mean in the reverse process as follows:

$$\begin{aligned} \mu_{t_{final}} = & \mu(x_t, t, M, F_1, F_2) \\ & + \gamma \sigma_\theta(x_t, t, M)^2 \log p_\psi(M|x_t) \end{aligned} \quad (9)$$

where  $\gamma$  controls the strength of the wire contrast. Typically,  $\gamma$  is randomly chosen from 0 (no segmentation-guided enhancement) to 15 to maintain visual realism.

## Experiments

**Datasets** Due to the absence of a publicly available cardiac fluoroscopy dataset, all data were collected from hospitals and manually annotated. Two datasets were used in this study: 1) a fully annotated fluoroscopy video (F-video) dataset consists of 400 2D+time fluoroscopy videos of 16 frames per video, from patients in cardiac interventional surgeries. Each video was normalized to an intensity range of -1 to 1. All frames are annotated with guidewire delineations by certified professionals. The frames were resized to  $512 \times 512$  pixels. 2) A partially annotated fluoroscopy image (P-image) dataset includes 409 2D+time fluoroscopy videos. Of these, 81 videos contain partial frames with guidewire annotations, and the rest videos have no annotations at all. Each video was normalized to an intensity range of -1 to 1. We collect 4000 annotated frames and 10000 unannotated frames, with all images resized to  $512 \times 512$  pixels, as the final dataset.

SF-VD was trained using both the P-image and F-video datasets, with both used for the scene-distribution model  $\theta_s$ , and the latter for the motion-distribution model  $\theta_m$ . Please refer to Appendix B for implementation details.

**Video Quality Analysis** We evaluated the quality of synthesized videos using three metrics: Fréchet Video Distance (FVD) (Unterthiner et al. 2018), Diversity Score (DS) (Pan et al. 2021), and Overfitting Score (OS) (Pan et al. 2021). FVD assesses the realism and diversity of the videos, DS measures the diversity among the videos, and OS evaluates whether the videos replicate training data or generate novel content.

**Wire Segmentation Evaluation** In addition, we evaluated the impact of fluoroscopy videos synthesized by SF-VD on a downstream wire segmentation task. Segmentation models were trained using 80% of the F-video dataset, augmented with 1200 synthesized videos featuring random wire contrasts generated by SF-VD. Performance was validated on another 10% of the F-video, and tested on the rest 10% of the F-video.

Metrics to evaluate segmentation performance include Dice Score (Dice) for region-overlap accuracy, Hausdorff Distance (HD) for boundary-deviation accuracy, Ground Truth to Result Error (G2RE) and Result to Ground Truth Error (R2GE) (Wu et al. 2008) for directional distance errors, and sensitivity and precision to address ambiguity in wire width and annotation variability. A more detailed explanation of these metrics is provided in Appendix B.

Various generic 3D segmentation networks, including V-Net (Milletari, Navab, and Ahmadi 2016), Dynvnet (Isensee et al. 2020), 3D Swin-Unetr (Hatamizadeh et al. 2021), and 3D Token-based MLP-Mixer (Pan et al. 2023a), were used in the downstream segmentation task. Models specifically designed for wire segmentation, such as PARN (Zhou et al. 2020), RTGN (Zhang et al. 2022), TTN (Zhang et al. 2021), and Ariadne+ (Caporali et al. 2022), were also included.

The efficacy of SF-VD augmentation was compared with other augmentation strategies, including geometric data augmentations (GA), pre-training methods such as Contrastive Learning (CL) (Chen et al. 2020) and Masked Volume Modeling (MVI) (Pathak et al. 2016), and synthesized videos from Video Diffusion Model (VD) (Ho et al. 2022) and Latent Video Diffusion Model (LVDM) (He et al. 2022).

Please refer to Appendix B for additional details on computing infrastructure, network implementation, and training.

## Results

Fig. 3 displays synthesized videos from the SF-VD, VD, and LVDM models. SF-VD generates realistic videos with diverse anatomical backgrounds and guidewire appearances. In contrast, VD, a 3D DDPM, results in noisy backgrounds and structural artifacts, highlighting the challenges of high-dimensional video generation without extensive datasets. LVDM-generated videos are visually realistic but exhibit limited variability, suggesting potential overfitting to a small annotated dataset. Conversely, SF-VD effectively produces varied videos by leveraging unannotated data. In addition, SF-VD generates realistic diaphragm motions, a capability that other methods fail to achieve. Fig. 4 shows synthetic videos with different wire contrasts, achieved through the segmentation-guided mechanism, which enhances the generalization capability and performance of the downstream segmentation network (Appendix D).

For quantitative evaluation, Table 1 shows that SF-VD achieves the highest FVD, indicating superior realism. Although it performs comparably to LVDM, SF-VD notably enhances video diversity, benefiting downstream segmentation tasks.

Table 2 shows that SF-VD significantly improves wire

Method	FVD ( $\downarrow$ )	DS ( $\downarrow$ )	OS ( $\downarrow$ )
SF-VD	<b>24.38</b>	<b>0.76±0.17</b>	<b>0.53±0.20</b>
VD	56.49	0.89±0.07	0.93±0.02
LVDM	27.55	0.91±0.05	0.92±0.04

Table 1: Comparison of the Fluoroscopy Video Quality Metrics for Different Methods. **BOLD** indicates the best result.

Method	Dynvnet	SwinUnetr	MLP-Mixer
No aug	0.52* / 28.38*	0.50* / 15.32*	0.52* / 11.85*
+GA	0.53* / 23.32*	0.51* / 15.50*	0.52* / 9.43*
+CL	0.54* / 22.39*	0.52* / 12.00*	0.54* / 10.50*
+MVI	0.56* / 23.92*	0.51* / 8.23*	0.53* / 10.38*
+VD	0.58 / 6.39*	0.58 / 7.51*	0.60 / 4.82*
+LVDM	0.55* / 4.83*	0.58* / 5.40*	0.55* / 4.32
<b>+SF-VD</b>	<b>0.61 / 3.93</b>	<b>0.61 / 4.23</b>	<b>0.61 / 4.22</b>

Method	PARN	RTGN	TTN
No aug	0.52* / 68.37	0.49* / 62.11*	0.52* / 44.17
+GA	0.54* / 29.38	0.50* / 45.02*	0.55* / 23.59
+CL	0.52* / 26.79	0.53* / 29.23*	0.50* / 16.40
+MVI	0.53* / 20.31	0.53* / 35.90*	0.48* / 8.39
+VD	0.56* / 12.57	0.62* / 5.90*	0.52* / 6.55
+LVDM	0.60* / <b>10.28</b>	0.61* / 2.77	0.52* / <b>6.21</b>
<b>+SF-VD</b>	<b>0.62 / 10.54</b>	<b>0.66 / 2.65</b>	<b>0.58 / 6.92</b>

Table 2: Comparison of data augmentation methods for wire segmentation using Dice ( $\uparrow$ ) and HD ( $\downarrow$ ): results for generic segmentation networks (upper table) and networks tailored for wire segmentation (lower table). **BOLD** indicates the best result. If SF-VD is the best, a paired t-test was conducted to compare it with the competing methods. \*indicates  $p < 0.05$ .

segmentation performance of various generic segmentation networks. It surpasses unaugmented models with significant improvements in both region-overlap accuracy (Dice) and boundary-deviation accuracy (HD). Compared to other augmentation methods (CL, MVI, GA, VD, and LVDM), SF-VD achieves superior or comparable Dice and HD. In Appendix C, SF-VD also shows better or comparable sensitivity, precision, and G2RE. In addition, SF-VD shows effectiveness in networks specifically designed for wire segmentation, generally outperforming competing augmentations in Dice, HD and other metrics (see Appendix C). This underscores SF-VD’s robust performance across various segmentation challenges.

**Ablation Study** We investigate the impact of Frame-Consistency (FC) and Segmentation-Guided (SG) strategies on the wire segmentation performance of V-Net. FC is designed to generate more realistic anatomical backgrounds and motion within videos, whereas SG focuses on modeling the diverse appearances of wires. Experiments without FC generate frames in a chronological order.

Results in Appendix D reveal that incorporating FC significantly improves metrics that measure region-overlap accuracy, including Dice, sensitivity, and precision. In contrast, distance-measuring metrics such as HD, G2RE, and R2GE show only minor changes with the addition of FC. Note that these distance-based metrics are more sensitive to small, isolated false-positive regions located far from the wires. Therefore, the likely explanation is that FC, by generating videos with more temporally coherent and realistic background, enables segmentation networks to learn features that better distinguish foreground and background motions, and thus reduce the extent of false positives and negatives. However, FC does not fully eliminate isolated false-positive regions. This could be because other structures in

cardiac fluoroscopy, such as ribs, likely exhibit motion patterns distinct from the background while also presenting sharp edges, which may be misclassified as wires.

SG, by contrast, greatly reduces HD and G2RE, effectively minimizing distance-based errors. It also improves R2GE, especially when FC is applied, although its influence on region-overlap metrics remains minor. This suggests that while SG does not substantially reduce false-positive or false-negative areas, it effectively eliminates isolated false-positive regions. A possible explanation is that by diversifying wire contrast in training data, SG helps segmentation networks learn more robust features for distinguishing wires from wire-like structures, such as rib edges. This finding highlights the importance of incorporating high variability in wire appearance within training data.

The simultaneous application of FC and SG yields the best results for Dice, HD, sensitivity, precision, and G2RE. This demonstrates that combining FC and SG strategies enhances both region-based and boundary-based metrics, achieving a well-rounded improvement.

**Limitation and Future Works** The proposed SF-VD is relatively inefficient due to its dual DDPM architecture. It requires 40 seconds to produce a 16-frame video on an A100 GPU, limiting both the number of videos and frames generated. Improving efficiency through a single model to handle all stages or by integrating high-efficiency DDPMs, such as the One-step Consistency Model (Song et al. 2023) or DiffuseVAE (Pandey et al. 2022), is crucial for efficient video generation within the proposed three-stage framework.

Additionally, current methods rely on existing wire mask sequences, limiting variability in wire motion. Future research will explore synthesizing both wire motion and background dynamics in the videos. For instance, synthesizing diverse wire mask sequences to guide video generation will allow for greater variability.

We also acknowledge that a limited number of training videos may affect the quality and variability of the generated outputs, potentially impacting downstream segmentation. Future work will analyze the relationship between training data size and segmentation performance to optimize data requirements.

Lastly, we aim to expand SF-VD to generate videos of other interventional devices, including catheters, stents, and catheter tips, to enhance segmentation and tracking for these devices.

## Conclusion

This study presents the SF-VD framework, a pioneering effort in medical video synthesis, aimed at creating a diverse library of fluoroscopy videos from existing guidewire annotations and augmenting the training data for deep learning models for wire segmentation. Our evaluations demonstrate the efficacy of these synthesized videos in improving deep learning-based wire segmentation. Future work will explore the potential of SF-VD framework for a broader range of interventional device-related tasks in fluoroscopy videos.

## References

- Avrahami, O.; Hayes, T.; Gafni, O.; Gupta, S.; Taigman, Y.; Parikh, D.; Lischinski, D.; Fried, O.; and Yin, X. 2023. Spatext: Spatio-textual representation for controllable image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18370–18380.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. 2023a. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023b. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22563–22575.
- Caporali, A.; Zanella, R.; Greogrio, D. D.; and Palli, G. 2022. Ariadne+: Deep Learning–Based Augmented Framework for the Instance Segmentation of Wires. *IEEE Transactions on Industrial Informatics*, 18(12): 8607–8617.
- Chen, H.; Zhang, Y.; Cun, X.; Xia, M.; Wang, X.; Weng, C.; and Shan, Y. 2024. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7310–7320.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion Models Beat GANs on Image Synthesis. *ArXiv*, abs/2105.05233.
- Dorjsembe, Z.; Pao, H.-K. K.; Odonchimed, S.; and Xiao, F. 2023. Conditional Diffusion Models for Semantic 3D Brain MRI Synthesis. *IEEE Journal of Biomedical and Health Informatics*, 28: 4084–4093.
- Du, Y.; Jiang, Y.; Tan, S.; Wu, X.; Dou, Q.; Li, Z.; Li, G.; and Wan, X. 2023. ArSDM: colonoscopy images synthesis with adaptive refinement semantic diffusion models. In *International conference on medical image computing and computer-assisted intervention*, 339–349. Springer.
- Gafni, O.; Polyak, A.; Ashual, O.; Sheynin, S.; Parikh, D.; and Taigman, Y. 2022. Make-a-scene: Scene-based text-to-image generation with human priors. In *European Conference on Computer Vision*, 89–106. Springer.
- Ge, S.; Nah, S.; Liu, G.; Poon, T.; Tao, A.; Catanzaro, B.; Jacobs, D.; Huang, J.-B.; Liu, M.-Y.; and Balaji, Y. 2023. Preserve your own correlation: A noise prior for video diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22930–22941.
- González, C.; Bravo-Sánchez, L.; and Arbelaez, P. 2020. ISINet: An Instance-Based Approach for Surgical Instrument Segmentation. In Martel, A. L.; Abolmaesumi, P.; Stoyanov, D.; Mateus, D.; Zuluaga, M. A.; Zhou, S. K.; Racoceanu, D.; and Joskowicz, L., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, 595–605. Cham: Springer International Publishing.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Gu, J.; Wang, S.; Zhao, H.; Lu, T.; Zhang, X.; Wu, Z.; Xu, S.; Zhang, W.; Jiang, Y.-G.; and Xu, H. 2023. Reuse and diffuse: Iterative denoising for text-to-video generation. *arXiv preprint arXiv:2309.03549*.
- Guo, Y.; Yang, C.; Rao, A.; Liang, Z.; Wang, Y.; Qiao, Y.; Agrawala, M.; Lin, D.; and Dai, B. 2023. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*.
- Hatamizadeh, A.; Nath, V.; Tang, Y.; Yang, D.; Roth, H. R.; and Xu, D. 2021. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In *International MICCAI Brainlesion Workshop*, 272–284. Springer.
- He, Y.; Yang, T.; Zhang, Y.; Shan, Y.; and Chen, Q. 2022. Latent video diffusion models for high-fidelity video generation with arbitrary lengths. *arXiv preprint arXiv:2211.13221*.
- Ho, J. 2022. Classifier-Free Diffusion Guidance. *ArXiv*, abs/2207.12598.
- Ho, J.; Chan, W.; Saharia, C.; Whang, J.; Gao, R.; Gritsenko, A.; Kingma, D. P.; Poole, B.; Norouzi, M.; Fleet, D. J.; et al. 2022. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022. Video Diffusion Models. *arXiv e-prints*, arXiv:2204.03458.
- Isensee, F.; Jaeger, P. F.; Kohl, S. A. A.; Petersen, J.; and Maier-Hein, K. 2020. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18: 203 – 211.
- Khachatryan, L.; Movsisyan, A.; Tadevosyan, V.; Henschel, R.; Wang, Z.; Navasardyan, S.; and Shi, H. 2023. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15954–15964.
- Konz, N.; Chen, Y.; Dong, H.; and Mazurowski, M. A. 2024. Anatomically-controllable medical image generation with segmentation-guided diffusion models. *arXiv preprint arXiv:2402.05210*.
- Lee, B. C.; Vaidya, K.; Jain, A. K.; and Chen, A. 2020. Guidewire Segmentation in 4D Ultrasound Sequences Using Recurrent Fully Convolutional Networks. In Hu, Y.; Licandro, R.; Noble, J. A.; Hutter, J.; Aylward, S.; Melbourne, A.; Abaci Turk, E.; and Torrents Barrena, J., eds., *Medical Ultrasound, and Preterm, Perinatal and Paediatric Image Analysis*, 55–65. Cham: Springer International Publishing. ISBN 978-3-030-60334-2.
- Liang, Z.; Li, Z.; Zhou, S.; Li, C.; and Loy, C. C. 2024. Control Color: Multimodal Diffusion-based Interactive Image Colorization. *arXiv preprint arXiv:2402.10855*.

- Liu, Y.; Lei, Y.; Fu, Y.; Wang, T.; Tang, X.; Jiang, X.; Curran, W. J.; Liu, T.; Patel, P.; and Yang, X. 2020. CT-based multi-organ segmentation using a 3D self-attention U-net network for pancreatic radiotherapy. *Medical physics*, 47(9): 4316–4324.
- Liu, Y.; Zhang, K.; Li, Y.; Yan, Z.; Gao, C.; Chen, R.; Yuan, Z.; Huang, Y.; Sun, H.; Gao, J.; et al. 2024. Sora: A review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177*.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11461–11471.
- Mao, J.; Wang, X.; and Aizawa, K. 2023. Guided image synthesis via initial image editing in diffusion model. In *Proceedings of the 31st ACM International Conference on Multimedia*, 5321–5329.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, 565–571. Ieee.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. In *International Conference on Machine Learning*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, 8162–8171. PMLR.
- Pan, S.; Chang, C.-W.; Wang, T.; Wynne, J.; Hu, M.; Lei, Y.; Liu, T.; Patel, P.; Roper, J.; and Yang, X. 2023a. Abdomen CT multi-organ segmentation using token-based MLP-Mixer. *Medical Physics*, 50(5): 3027–3038.
- Pan, S.; Flores, J.; Lin, C. T.; Stayman, J. W.; and Gang, G. J. 2021. Generative adversarial networks and radiomics supervision for lung lesion synthesis. In *Medical Imaging 2021: Physics of Medical Imaging*, volume 11595, 167–172. SPIE.
- Pan, S.; Wang, T.; Qiu, R. L. J.; Axente, M.; Chang, C.-W.; Peng, J.; Patel, A. B.; Shelton, J.; Patel, S. A.; Roper, J.; and Yang, X. 2023b. 2D medical image synthesis using transformer-based denoising diffusion probabilistic model. *Physics in Medicine & Biology*, 68(10): 105004.
- Pandey, K.; Mukherjee, A.; Rai, P.; and Kumar, A. 2022. Diffusevae: Efficient, controllable and high-fidelity generation from low-dimensional latents. *arXiv preprint arXiv:2201.00308*.
- Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2536–2544.
- Prasad, M.; Cassar, A.; Fetterly, K. A.; Bell, M.; Theessen, H.; Ecabert, O.; Bresnahan, J. F.; and Lerman, A. 2016. Co-registration of angiography and intravascular ultrasound images through image-based device tracking. *Catheterization and Cardiovascular Interventions*, 88(7): 1077–1082.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency models. *arXiv preprint arXiv:2303.01469*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Tang, Y.; Yang, D.; Li, W.; Roth, H. R.; Landman, B.; Xu, D.; Nath, V.; and Hatamizadeh, A. 2022. Self-supervised pre-training of swin transformers for 3d medical image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20730–20740.
- Unterthiner, T.; Van Steenkiste, S.; Kurach, K.; Marinier, R.; Michalski, M.; and Gelly, S. 2018. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*.
- Vöth, T.; Koenig, T.; Eulig, E.; Knaup, M.; Wiesmann, V.; Hörndler, K.; and Kachelrieß, M. 2023. Real-time 3D reconstruction of guidewires and stents using two update X-ray projections in a rotating imaging setup. *Medical Physics*, 50(9): 5312–5330.
- Wang, W.; Bao, J.; Zhou, W.; Chen, D.; Chen, D.; Yuan, L.; and Li, H. 2022. Semantic image synthesis via diffusion models. *arXiv preprint arXiv:2207.00050*.
- Wu, C.; Liang, J.; Ji, L.; Yang, F.; Fang, Y.; Jiang, D.; and Duan, N. 2022. Nüwa: Visual synthesis pre-training for neural visual world creation. In *European conference on computer vision*, 720–736. Springer.
- Wu, Z.; Sun, J.; Liu, Q.; and Zhang, G. 2008. A Novel Calibration Method for Large-Scale Stereo Vision Sensor based on One-Dimensional Target. 7129.
- Yang, R.; Srivastava, P.; and Mandt, S. 2023. Diffusion probabilistic modeling for video generation. *Entropy*, 25(10): 1469.
- Yu, X.; Li, G.; Lou, W.; Liu, S.; Wan, X.; Chen, Y.; and Li, H. 2023. Diffusion-based data augmentation for nuclei image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 592–602. Springer.
- Zhang, B.; Bui, M.; Wang, C.; Bourier, F.; Schunkert, H.; and Navab, N. 2022. Real-time guidewire tracking and segmentation in intraoperative x-ray. In *Medical Imaging 2022: Image-Guided Procedures, Robotic Interventions, and Modeling*, volume 12034, 283–290. SPIE.
- Zhang, G.; Wong, H.-C.; Wang, C.; Zhu, J.; Lu, L.; and Teng, G. 2021. A Temporary Transformer Network for Guide-Wire Segmentation. In *2021 14th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, 1–5.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3813–3824.
- Zhou, Y.-J.; Xie, X.-L.; Zhou, X.-H.; Liu, S.-Q.; Bian, G.-B.; and Hou, Z.-G. 2020. Pyramid attention recurrent networks for real-time guidewire segmentation and tracking in intraoperative X-ray fluoroscopy. *Computerized Medical Imaging and Graphics*, 83: 101734.