

SG2VID: Scene Graphs Enable Fine-Grained Control for Video Synthesis

Ssharvien Kumar Sivakumar¹[0009–0009–0383–9462] Yannik Frisch^{1,2},
Ghazal Ghazaei³, and Anirban Mukhopadhyay¹

¹ Technical University Darmstadt, Darmstadt, Germany
`ssharvien_kumar.sivakumar@tu-darmstadt.de`

² Universitätsmedizin der Johannes Gutenberg-Universität Mainz, Mainz, Germany

³ Carl Zeiss AG, Munich, Germany

Abstract. Surgical simulation plays a pivotal role in training novice surgeons, accelerating their learning curve and reducing intra-operative errors. However, conventional simulation tools fall short in providing the necessary photorealism and the variability of human anatomy. In response, current methods are shifting towards generative model-based simulators. Yet, these approaches primarily focus on using increasingly complex conditioning for precise synthesis while neglecting the fine-grained human control aspect. To address this gap, we introduce SG2VID, the first diffusion-based video model that leverages Scene Graphs for both precise video synthesis and fine-grained human control. We demonstrate SG2VID’s capabilities across three public datasets featuring cataract and cholecystectomy surgery. While SG2VID outperforms previous methods both qualitatively and quantitatively, it also enables precise synthesis, providing accurate control over tool and anatomy’s size and movement, entrance of new tools, as well as the overall scene layout. We qualitatively motivate how SG2VID can be used for generative augmentation and present an experiment demonstrating its ability to improve a downstream phase detection task when the training set is extended with our synthetic videos. Finally, to showcase SG2VID’s ability to retain human control, we interact with the Scene Graphs to generate new video samples depicting major yet rare intra-operative irregularities.

Keywords: Simulation · Scene Graph Conditioning · Video Generation

1 Introduction

The Halsted model of "See One, Do One, Teach One" places patients at the centre of novice surgeons’ training [32]. This often results in suboptimal patient outcomes due to the limited abilities of novice surgeons to handle intra-operative complications [15]. Integrating simulation tools into surgical training significantly accelerates their learning curve and reduces the likelihood of procedural mistakes [20]. However, existing simulation methods that come in the form of generic phantoms and computer-rendered graphics lack visual realism

as the size of scene entities, or the entry of new entities, e.g. additional surgical instruments. This is shown in our qualitative comparison in the supplementary.

We propose **SG2VID, leveraging SG conditioning for realistic surgical video synthesis with precise and fine-grained human control**. The SGs containing spatial and temporal information, as depicted in Figure 1, are used to pre-train two separate graph encoders, capturing local and global features from the surgical scene into an intermediate representation. This representation is subsequently used to condition a video diffusion model [23]. The first-frame of any surgical sequence can optionally be utilised as an additional conditioning signal. This is relevant for surgical simulation and analogous to generative editing, where a patient-specific reference image can be used for personalised simulation. In summary, our contributions are as follows: **(I)** SG2VID is the **first diffusion-based SG-to-Video model** that enables fine-grained control for video synthesis. The effectiveness of SG2VID is demonstrated across three cataract and cholecystectomy surgery datasets. **(II)** We showcase SG2VID’s generalisation capabilities by generating sequences with precise control from first-frame images of video domains outside of the training data. We also demonstrate how rare samples of intra-operative irregularities can be generated by interactively manipulating the graph, emphasising the human controllability aspect. **(III)** We highlight our method’s potential for generative augmentation by improving a downstream model for surgical phase recognition with the dataset extended with SG2VID’s synthetic data.

2 Method

We define a Scene Graph as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v^1, v^2, \dots, v^k\}$ represents a set of k nodes. The nodes are connected by a set of undirected edges \mathcal{E} . The Scene Graph \mathcal{G} is constructed using segmentation mask m belonging to an image x . Given that the public surgical datasets we utilise provide segmentation annotations for only a small set of videos on limited frames, we employ SASVi [25] to generate segmentation masks for each frame across all videos. Each node corresponds to a connected component within the mask m , such as a tool or anatomical structure. In line with SurGrID [6], each node $v^j \in \mathbb{R}^{d+7}$ is encoded with a d -dim class vector, a 2-dim centroid position, and 2-dim spatial spread representing the height and width of the component. We further extend this representation to include the 2-dim average optical flow direction and the 1-dim average depth. This additional information assists in modelling the motion trajectory by capturing temporal dynamics within the surgical scene. To compute the average flow and depth values, we leverage pre-trained RAFT [28] and MiDaS [22] models averaged over each component’s region. Finally, edges \mathcal{E} in the graph \mathcal{G} connect nodes \mathcal{V} whose corresponding connected components are spatially adjacent in the mask m .

2.1 Pre-Training Graph Encoder

To enable the graph conditioning of our video diffusion model, we employ two graph encoders, E_G^{glob} and E_G^{loc} , which share the same architecture but are trained with different learning objectives. E_G^{glob} captures global scene dynamics, including the overall layout, interactions, and motion patterns, while E_G^{loc} focuses on encoding local information such as the texture and fine-grained details of the components. The graph embeddings from both encoders are concatenated $z_G = \text{concat}(z_G^{glob}, z_G^{loc})$ and used to condition the video diffusion model. Each of the graph encoders is constructed using a series of Graph Attention Networks (GATv2) layers [2], where each node updates its representation by attending to its neighbours using its own representation as the query [2]. A graph-level representation is obtained through mean pooling of the node-level representations.

Given a video sequence consisting of image frames $x_{1:n} = \{x_1, x_2, \dots, x_n\}$ and the corresponding graphs $\mathcal{G}_{1:n} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n\}$, we randomly select a component, such as a tool or anatomical structure, and mask it across the entire sequence $\tilde{x}_{1:n} = \{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n\}$. The graph encoder E_G^{loc} , together with an auxiliary transformer-based decoder d_θ , is trained to predict the complete sequence $x_{1:n}$ given the masked sequence $\tilde{x}_{1:n}$ and the graph $\mathcal{G}_{1:n}$, by minimizing the reconstruction loss in Equation 1. This operation takes place in the embedding space, where $z_x = E_x(x)$ and $z_{\tilde{x}} = E_x(\tilde{x})$, but for simplicity, we keep the notations in the equations in the original form.

$$\mathcal{L}_{loc} = \mathbb{E}_{x_{1:n}, \mathcal{G}_{1:n}} \|x_{1:n} - d_\theta(\tilde{x}_{1:n}, \mathcal{G}_{1:n})\|_2^2 \quad (1)$$

$$\mathcal{L}_{glob} = \mathbb{E}_{\mathcal{G}_{1:n}, m_{1:n}^+, m_{1:n}^-} \left[-\log \frac{\exp(\mathcal{G}_{1:n} \cdot m_{1:n}^+)}{\exp(\mathcal{G}_{1:n} \cdot m_{1:n}^+) + \sum_i \exp(\mathcal{G}_{1:n} \cdot m_{1:n}^-)} \right] \quad (2)$$

In parallel, we train E_G^{glob} using the contrastive learning objective in Equation 2, aligning the embedding spaces of the graphs $\mathcal{G}_{1:n}$ and segmentation masks $m_{1:n}$. This optimisation is also performed in embedding space, where segmentation masks are encoded using $z_m = E_m(m)$, but the notation is kept in its original form for simplicity. By minimizing Equation 2, E_G^{glob} is forced to bring compliant target masks $m_{1:n}^+$ closer to $\mathcal{G}_{1:n}$, while pushing the negative target masks $m_{1:n}^-$ further away in the embedding space. By operating in the segmentation space, the model only focuses on capturing the overall layout and ignores textural information, as such details are absent in the masks.

2.2 Graph-Image to Video Diffusion Model

Diffusion Models (DMs) [10] employ a parameterized model ϵ_θ , typically a 2D U-Net, that learns to model the data distribution by approximating the reversal of gradual noising $p(x_{t-1}|x_t, c)$ by minimizing the error between predicted and target noise: $\min_\theta \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t, c)\|^2]$. Here, c denotes the conditioning signal, which in this work comprises the graph embeddings $z_{\mathcal{G}_{1:n}}$ and its

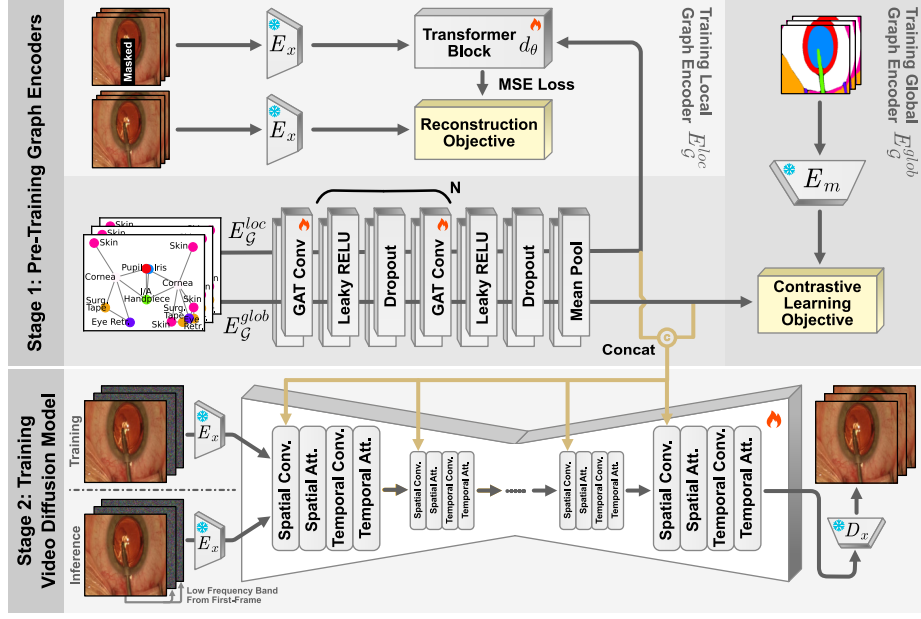


Fig. 2. Training Workflow of the Graph Encoders and Video Diffusion Model.

first-frame x_1 . For computational efficiency, Rombach et al. [24] propose to perform the forward and backward process, $p(z_{t-1}|z_t, c)$, in the latent space. So to apply the first-frame conditioning, we encode the first-frame $z_1 = E_x(x_1)$ and replace it with the first-frame noise ϵ^1 . The model input is then constructed as $\hat{\epsilon} = \{z^1, \epsilon^2, \epsilon^3, \dots, \epsilon^N\}$. To incorporate the graph conditioning signal, the graph embeddings $z_{G_{1:n}}$ are concatenated with the timestep embedding before being passed to the U-Net blocks. Finally, the sample output from the model can be obtained using a pre-trained decoder $\hat{x} = D_x(\hat{z})$.

For video generation using diffusion models [11], the parameterised model is adapted with a spatial-temporal factorised 3D U-Net. A temporal layer, consisting of temporal convolutions and temporal attention, is introduced after each spatial layer. The interleaving of the temporal layers with the existing spatial layers brings temporal awareness to the model. To further improve motion consistency in generated videos and ensure visual coherence from the first-frame throughout the sequence, we incorporate strategies from ConsistI2V [23]. Specifically, a local window of first-frame features is incorporated into the temporal layers to augment their attention operations. Common noise schedules have been shown to introduce information leakage about diffusion noise during training, leading to a discrepancy between the initialisation of noise for training and sampling [18]. To address this issue, ConsistI2V [23] integrates the low-frequency components of the first-frame image into the initial noise during inference.

3 Experiments and Results

We train and validate our model on three publicly available surgical video datasets featuring real-world cataract and cholecystectomy surgeries: Cataract-1k [7], CATARACTS [1], and Cholec80 [29]. These datasets vary significantly in size, containing 1000, 50, and 80 videos, respectively, and also exhibit substantial differences in the average video length. To standardise the training, videos are downsampled to 4 fps, and video sequences of 16 frames are extracted with a 75% overlap between neighbouring sequences. The frames are downsampled to a resolution of 128×128. The train/validation/test split was assigned on a video-wise basis: 80/10/10% for Cholec80 and Cataract-1k, and 50/6/44% for CATARACTS. The graph encoders are trained on a single A40 GPU, while the video diffusion model is trained on four A40 GPUs, though training with a single A40 GPU using gradient accumulation remains feasible. Upon acceptance, we will release our SG representations, model weights, and code at github.com/MECLabTUDA/SG2VID.

Table 1. Quantitative Assessment of Video Synthesis Quality.

Method	Cataract-1k [7]			CATARACTS [1]			Cholec80 [29]		
	FVD↓	FID↓	LPIPS↑	FVD↓	FID↓	LPIPS↑	FVD↓	FID↓	LPIPS↑
StyleGAN-V [26]	442.6	118.2	0.286	618.7	94.8	0.382	1544.1	200.5	0.378
Endora [16]	265.9	30.2	0.377	649.5	45.9	0.454	533.8	47.0	0.525
MedSora [31]	901.8	137.6	0.324	952.1	112.1	0.403	1297.5	153.6	0.406
*LVDM [9]	1656.6	186.0	0.534	1178.9	113.7	0.559	1507.0	110.8	0.668
MOFA [21]	722.2	89.9	0.361	713.4	88.3	0.460	651.1	72.4	0.506
SG2VID (Ours)	77.0	15.5	0.397	523.8	40.9	0.444	457.3	16.4	0.532
SG2VID-XIMG (Ours)	278.4	33.3	0.409	535.7	39.8	0.465	560.1	25.1	0.533

Method	Cataract-1k [7]		CATARACTS [1]		Cholec80 [29]	
	BB IoU↑	F1↑	BB IoU↑	F1↑	BB IoU↑	F1↑
*LVDM [9]	0.193	0.164	0.225	0.149	0.42	0.28
MOFA [21]	0.456	0.449	0.425	0.275	0.599	0.45
SG2VID (Ours)	0.624	0.634	0.493	0.379	0.623	0.476
SG2VID-XIMG (Ours)	0.440	0.460	0.389	0.318	0.530	0.389
Mask R-CNN on Annotated Subset	0.731	0.745	0.636	0.585	0.887	0.957

Quantitative Assessment: To quantitatively assess the quality and diversity of the generated sequences, we utilise a combination of FID, FVD [30] and LPIPS diversity score [34], with results shown in Table 1 (*top*). Additionally, to assess how well the generated sequences align with the SG conditioning, we utilise a Mask R-CNN [8] model. The model is pre-trained on the datasets’ subsets with segmentation annotations. Doing so allows us to compare the model’s predictions between the real and generated sequences, evaluating F1 scores (for object classes) and bounding box IoU (for object localisation). The results of this

evaluation are reported in Table 1 (*bottom*). The Mask R-CNN’s performance on the annotated subset serves as an upper bound here.

We compare our method against various approaches with publicly available repositories. StyleGAN-V [26], Endora [16], and MedSora [31] are unconditional video generation models. We adapt LVDM [9] for text-based conditioning and use textual descriptions from SG triplets, as in VISAGE [33]. For trajectory-based conditioning, MOFA [21] employs a combination of first-frame and motion trajectories, obtained by applying sparse motion sampling to optical flow, similar to the strategy of SurgSora [4]. As shown in Table 1, SG2VID surpasses all the baselines in terms of visual fidelity, while adhering to the conditioning signal. SG2VID-XIMG is our model trained using graphs alone without the first-frame.

Ablation: We perform an ablation study on Cataract-1k to highlight the importance of jointly modelling local and global cues. SG2VID’s local-only variant scores an FVD of 80.6, FID of 16.4, and LPIPS of 0.396, while the global-only variant scores 79.8, 15.9, and 0.397, both underperforming SG2VID.




Fig. 3. SG-to-Video Generation. More samples are in the supplementary material.

Qualitative Assessment: Figure 3 presents qualitative results of SG2VID and SG2VID-XIMG, demonstrating how SG conditioning enables precise synthesis. In the synthesised sequences, the movement and size of the tools and anatomies, as well as the overall layout, closely follow the ground truth. SG2VID also has no difficulty synthesising the entry of a brand-new tool absent in the first-frame, as well as the exit and reentry of existing tools. Without the first-frame conditioning, SG2VID-XIMG still accurately captures tool and anatomy movements, but video-specific features absent from the SG, such as eye colour,

appear random, akin to generative augmentation. We provide additional samples of our synthetic videos in the supplementary material, including qualitative comparisons to related work.

Method Controllability and Generalisation: To highlight human controllability and generalisation to uncommon cases, we simulate sudden pupil contraction, known as miosis, a rare intra-operative irregularity described in Cataract-1k [7]. This phenomenon risks significant tissue injury, especially when a tool is deeply inserted. We select a set of SGs and reduce the pupil size in the final SG. The pupil size values for the intermediate SGs are simply interpolated between the initial and final SG. Figure 4 (*top*) shows synthesised sequences with sudden pupil contractions. To demonstrate generalisation, we generate sequences from entirely unseen eye images to recreate unseen domain style transfer. We take eye images from the OphNet dataset [13] and condition on SGs from Cataract-1k to synthesise new video sequences, shown in Figure 4 (*bottom*). We also provide more samples for both cases in the supplementary material.




Fig. 4. Synthesis of Pupil Contraction Irregularities and Generalisation on Unseen Images from the OphNet Dataset. More samples are in the supplementary material.

Generative Augmentation on Surgical Phase Recognition: The qualitative results in Figure 3 motivate the use of SG2VID as a generative augmenter. We select surgical phase classification on CATARACTS [1] as a downstream task and train on our generatively extended dataset. To extend the training set, we use first-frame from one video and SGs from another (mix-and-match) within the training set of CATARACTS [1] to generate new video sequences. Since our model allows for precise control over most components during the synthesis process, this would be comparable to sequence-to-sequence translation. Given that our generated sequences are only 16 frames long, we generate each subsequent sequence using the last frame of the previous one, producing complete surgical videos from start to finish. Phase annotations for synthesised videos are the same as the target videos from which we take the graph. We train MS-TCN++ [17] model using features from DINO [3] on both the original and the

extended datasets containing our generated videos. The macro F1-Score and Accuracy when trained on the normal training set are 0.794 and 0.793, respectively. Extending the training set with our synthesised videos improves the downstream phase recognition, yielding macro F1 and Accuracy scores of 0.805 and 0.816.

4 Conclusion

We present the **first diffusion-based SG-to-Video model (SG2VID)**, enabling fine-grained video synthesis through human-interactable conditioning. Our method works across multiple surgical domains, demonstrating both qualitative and quantitative superiority. SG2VID provides precise control over the movement and size of tools and anatomical structures, as well as new tool entry and exit. To showcase its human controllability, we generate rare pupil collapse irregularities in cataract surgery by interactively manipulating the SGs. Our results highlight SG2VID’s potential as a generative augmentor, improving downstream phase recognition with a generatively extended training set.

Acknowledgments. The work is partially funded by the German Federal Ministry of Education and Research as part of the Software Campus (research grant 01|S23067).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Al Hajj, H., Lamard, M., Conze, P.H., Roychowdhury, S., Hu, X., Maršalkaitė, G., Zisimopoulos, O., Dedmari, M.A., Zhao, F., Prellberg, J., et al.: Cataracts: Challenge on automatic tool annotation for cataract surgery. *MedIA* **52**, 24–41 (2019)
2. Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? *arXiv:2105.14491* (2021)
3. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *ICCV*. pp. 9650–9660 (2021)
4. Chen, T., Yang, S., Wang, J., Bai, L., Ren, H., Zhou, L.: Surgsora: Decoupled rgb-d-flow diffusion model for controllable surgical video generation. *arXiv:2412.14018* (2024)
5. Cho, J., Schmidgall, S., Zakka, C., Mathur, M., Kaur, D., Shad, R., Hiesinger, W.: Surgen: Text-guided diffusion model for surgical video generation. *arXiv:2408.14028* (2024)
6. Frisch, Y., Sivakumar, S.K., Köksal, Ç., Böhm, E., Wagner, F., Gericke, A., Ghazaei, G., Mukhopadhyay, A.: Surgrid: Controllable surgical simulation via scene graph to image diffusion. *arXiv:2502.07945* (2025)
7. Ghamsarian, N., El-Shabrawi, Y., Nasirihaghighi, S., Putzgruber-Adamitsch, D., Zinkernagel, M., Wolf, S., Schoeffmann, K., Sznitman, R.: Cataract-1k: Cataract surgery dataset for scene segmentation, phase recognition, and irregularity detection. *arXiv:2312.06295* (2023)

8. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV. pp. 2961–2969 (2017)
9. He, Y., Yang, T., Zhang, Y., Shan, Y., Chen, Q.: Latent video diffusion models for high-fidelity long video generation. arXiv:2211.13221 (2022)
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *NeurIPS* **33**, 6840–6851 (2020)
11. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. *NeurIPS* **35**, 8633–8646 (2022)
12. Holm, F., Ghazaei, G., Czempiel, T., Özsoy, E., Saur, S., Navab, N.: Dynamic scene graph representation for surgical video. In: ICCV Workshop. pp. 81–87 (2023)
13. Hu, M., Xia, P., Wang, L., Yan, S., Tang, F., Xu, Z., Luo, Y., Song, K., Leitner, J., Cheng, X., et al.: Ophnet: A large-scale video benchmark for ophthalmic surgical workflow understanding. In: ECCV. pp. 481–500. Springer (2024)
14. Iliash, I., Allmendinger, S., Meissen, F., Köhl, N., Rückert, D.: Interactive generation of laparoscopic videos with diffusion models. In: MICCAI Workshop on Deep Generative Models. pp. 109–118. Springer (2024)
15. Kwong, A., Law, S.K., Kule, R.R., Nouri-Mahdavi, K., Coleman, A.L., Caprioli, J., Giacon, J.A.: Long-term outcomes of resident-versus attending-performed primary trabeculectomy with mitomycin c in a united states residency program. *American journal of ophthalmology* **157**(6), 1190–1201 (2014)
16. Li, C., Liu, H., Liu, Y., Feng, B.Y., Li, W., Liu, X., Chen, Z., Shao, J., Yuan, Y.: Endora: Video generation models as endoscopy simulators. In: MICCAI. pp. 230–240. Springer (2024)
17. Li, S., Farha, Y.A., Liu, Y., Cheng, M.M., Gall, J.: Ms-tcn++: Multi-stage temporal convolutional network for action segmentation. *IEEE PAMI* **45**(6), 6647–6658 (2020)
18. Lin, S., Liu, B., Li, J., Yang, X.: Common diffusion noise schedules and sample steps are flawed. In: WACV. pp. 5404–5411 (2024)
19. Murali, A., Alapatt, D., Mascagni, P., Vardazaryan, A., Garcia, A., Okamoto, N., Mutter, D., Padoy, N.: Latent graph representations for critical view of safety assessment. *IEEE TMI* (2023)
20. Nair, A.G., Ahiwalay, C., Bacchav, A.E., Sheth, T., Lansingh, V.C., Vedula, S.S., Bhatt, V., Reddy, J.C., Vadavalli, P.K., Praveen, S., et al.: Effectiveness of simulation-based training for manual small incision cataract surgery among novice surgeons: a randomized controlled trial. *Scientific reports* **11**(1), 10945 (2021)
21. Niu, M., Cun, X., Wang, X., Zhang, Y., Shan, Y., Zheng, Y.: Mofa-video: Controllable image animation via generative motion field adaptations in frozen image-to-video diffusion model. In: ECCV. pp. 111–128. Springer (2024)
22. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE PAMI* **44**(3), 1623–1637 (2020)
23. Ren, W., Yang, H., Zhang, G., Wei, C., Du, X., Huang, W., Chen, W.: Consisti2v: Enhancing visual consistency for image-to-video generation. arXiv:2402.04324 (2024)
24. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
25. Sivakumar, S.K., Frisch, Y., Ranem, A., Mukhopadhyay, A.: Sasvi-segment any surgical video. arXiv:2502.09653 (2025)
26. Skorokhodov, I., Tulyakov, S., Elhoseiny, M.: Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In: CVPR. pp. 3626–3636 (2022)

27. Sun, W., You, X., Zheng, R., Yuan, Z., Li, X., He, L., Li, Q., Sun, L.: Bora: Biomedical generalist video generation model. arXiv:2407.08944 (2024)
28. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: ECCV. pp. 402–419. Springer (2020)
29. Twinanda, A.P., Shehata, S., Mutter, D., Marescaux, J., De Mathelin, M., Padoy, N.: Endonet: a deep architecture for recognition tasks on laparoscopic videos. IEEE TMI **36**(1), 86–97 (2016)
30. Unterthiner, T., Van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Towards accurate generative models of video: A new metric & challenges. arXiv:1812.01717 (2018)
31. Wang, Z., Zhang, L., Wang, L., Zhu, M., Zhang, Z.: Optical flow representation alignment mamba diffusion model for medical video generation. arXiv:2411.01647 (2024)
32. Wood, T.C., Maqsood, S., Sancha, W., Saunders, A., Lockington, D., Nanavaty, M.A., Rajak, S.: Principles of simulation and their role in enhancing cataract surgery training. Eye **36**(8), 1529–1531 (2022)
33. Yeganeh, Y., Lazuardi, R., Shamseddin, A., Dari, E., Thirani, Y., Navab, N., Farshad, A.: Visage: Video synthesis using action graphs for surgery. In: MICCAI. pp. 146–156. Springer (2024)
34. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595 (2018)
35. Zhou, X., Huang, Y., Xue, W., Dou, H., Cheng, J., Zhou, H., Ni, D.: Heartbeat: Towards controllable echocardiography video synthesis with multimodal conditions-guided diffusion models. In: MICCAI. pp. 361–371. Springer (2024)