



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

DeGenseGS: Geometrically and Semantically Decoupled Surgical Scene Understanding in 4D Gaussian Splatting

Yimo Wang^{1,3*}, Bin Kang^{2*}, Shuojue Yang³, and Yueming Jin^{3**}

¹ School of Automation, Southeast University, Nanjing 210096, China
yimowang@seu.edu.cn

² School of Internet of Things, Nanjing University of Posts and Telecommunications, Nanjing, China
kangb@njupt.edu.cn

³ Department of Biomedical Engineering, National University of Singapore, Singapore
{yimowang,s.yang}@u.nus.edu, ymjn@nus.edu.sg

Abstract. Real-time, text-promptable 4D reconstruction is indispensable for autonomous surgical interaction. Severe misalignment between semantic meaning and physical anatomy still persists, largely because existing solutions integrate Vision-Language Models into deformable fields via a rigid coupling scheme that tightly binds semantic features to geometric warping. In this paper, we propose DeGenseGS, Geometrically and Semantically Decoupled Surgical Scene Understanding in 4D Gaussian Splatting, a novel framework that independently models semantic evolution and geometric deformation. Specifically, we propose a HexPlane-based spatiotemporal entanglement module that uses shared kinematic latents to synchronize semantic mutations with scene dynamics, while explicitly disentangling semantic updates from geometric deformation. To further ensure robustness against reconstruction artifacts, we devise a Rasterization-Native Semantic Extraction mechanism that infers semantics from topologically continuous feature maps. Additionally, we incorporate an angular-aligned optimization strategy that conforms to the native hyperspherical latent space, thereby preventing semantic distortion. Extensive evaluations on the CholecSeg8k and EndoVis18 datasets demonstrate that DeGenseGS achieves state-of-the-art performance. Our framework yields enhanced geometric completeness and robust semantic-anatomic alignment, enabling spatially continuous segmentation despite drastic tissue deformation and topological transitions.

Keywords: 4D Gaussian Splatting · Surgical Scene Understanding · Vision-Language Models · Semantic-Geometric Decoupling.

* Yimo Wang and Bin Kang contributed equally to this work.

** Corresponding author.

1 Introduction

In the progression towards autonomous Robotic-Assisted Minimally Invasive Surgery (RAMIS), surgical perception systems must evolve from passive geometric observers to active cognitive agents. While high-fidelity 4D reconstruction provides essential spatial mapping, it remains semantically agnostic[9]. To enable advanced downstream capabilities, such as safety-critical zone alerting and fine-grained instrument tracking, the spatial reconstruction must be augmented with text-promptable semantic understanding. This requires a framework capable of simultaneously tracking dynamic anatomy in real-time and assigning queryable semantic identities to the reconstructed environment, thereby bridging the gap between low-level geometric sensing and high-level clinical reasoning.

Recent advancements have transitioned surgical scene reconstruction from implicit neural representations [21, 24] to explicit 4D Gaussian Splatting (4DGS) [7, 16, 25, 2], achieving remarkable rendering efficiency and visual fidelity. Concurrently, surgical Vision-Language Models (VLMs)[20, 12, 10, 17] have enabled sophisticated text-promptable 2D scene interpretation [15]. Despite these parallel successes, a critical dimensionality gap persists: dynamic 4D reconstructions inherently lack semantic awareness, whereas 2D VLMs lack essential spatial grounding [9]. While recent pioneering efforts [18, 8, 6] attempt to bridge this gap by lifting 2D semantic priors into dynamic volumetric fields, establishing robust semantic-geometric alignment remains an open challenge. Specifically, the severe occlusions, complex tissue deformations, and rapid topological changes characteristic of endoscopic environments [19, 14] expose the fragility of these early rigidly coupled representations. This inherent limitation necessitates a specialized framework to explicitly decouple spatial deformation from semantic evolution.

To populate an independent semantic field, extracting robust 2D priors is a prerequisite. While SAM-based models yield sharp boundaries, their features lack explicit semantic representation [28, 29]. Accordingly, incorporating Vision-Language Models (VLMs) is indispensable for text-promptable surgical scene understanding. The key challenge lies in simultaneously modeling semantic evolution alongside geometric deformation. Current 3DGS baselines cast these two factors into a single deformation framework, leveraging the Flexible Deformation Model (FDM) that couples spatial warping with semantic refinement in a shared pipeline [25, 2, 22]. This coupled design suffers from the gradient sensitivity issue [26] in surgical events such as surgical cauterization, where tissue appearance mutates significantly but the physical geometry remains largely static. Specifically, we observe that the network gradients of FDM are far more sensitive to semantic shifts than to subtle geometric changes [30]. As semantic and geometric representations share the same FDM weight, gradients dominated by semantic changes will become excessively large in traditional solutions, inadvertently leading to spurious geometric warping during the surgical cauterization task. As shown in Fig. 1, there exist misaligned representations for semantic evolution and geometric deformation in FDM. These erroneous deformations propagate to the




Fig. 1. Illustration of spatiotemporal feature misalignment in existing coupled 4D Gaussian Splatting frameworks. **(A) Physical Reality:** During topology-altering events such as cauterization, tissue appearance mutates significantly while the underlying geometry remains largely stable. **(B) Limitation of existing Coupled Frameworks:** Due to the gradient sensitivity variation, the coupled framework leads to unfounded geometric warping (blue arrows) to accommodate the new semantic identity.

semantic mask, eventually undermining the geometric and semantic consistency of the reconstructed scene.

In this paper, we propose DeGenseGS, namely Geometrically and Semantically Decoupled Surgical Scene Understanding in 4D Gaussian Splatting, which integrates VLM in 4D Gaussian Splatting for separately model semantic evolution and geometric deformation. First, at the spatiotemporal level, a HexPlane-based module is proposed to extract shared kinematic latents. These latents are routed to two decoupled decoding branches to independently model geometric deformation and semantic evolution. This design mitigates gradient interference, preventing semantic losses from dominating the optimization of geometric parameters. As a result, tissue appearance can evolve freely without inducing spurious spatial displacements. Second, to reduce reconstruction artifacts, we propose a Rasterization-Native Semantic Extraction mechanism that decouples semantic features from deformed 4D Gaussian kernels. It performs graph-based analysis on topologically continuous 2D feature maps and leverages RGB renderings as structural guidance to correct local VLM inconsistencies, producing robust, hole-free segmentation masks. Finally, we introduce a manifold-aligned optimization paradigm for robust semantic grounding, where the semantic distillation is formulated as a metric learning problem. The contributions of this paper are as follows:

- We propose DeGenseGS, the first text-promptable 4DGS framework that explicitly decouples spatial deformation from semantic evolution, enabling fine-grained surgical interaction.

- We introduce a Rasterization-Native Semantic Extraction mechanism that leverages RGB structural guidance to rectify VLM inconsistencies and ensure hole-free, precise segmentation boundaries.
- We propose a Manifold-Aligned Optimization strategy to regularize the VLM latent space via angular distillation, effectively mitigating feature collapse in dynamic surgical environments.
- State-of-the-art performance is validated on comprehensive datasets, demonstrating superior semantic-geometric alignment capability.

2 Methodology

DeGenseGS is a 4D Gaussian Splatting framework designed for robust surgical scene understanding (Fig. 2). It explicitly decouples the representation of semantic evolution from spatial deformation using a kinematics-conditioned latent disentanglement mechanism. And it also recovers precise boundaries via a rasterization-native semantic extraction.

2.1 Preliminaries

3D Gaussian Splatting (3DGS) represents scenes using anisotropic 3D Gaussians, defined by center $\boldsymbol{\mu} \in \mathbb{R}^3$, covariance $\boldsymbol{\Sigma}$, opacity σ , and color c . To render a novel view, 3D Gaussians are projected into 2D splats and alpha-blended in depth order: $\hat{\mathbf{C}}(\mathbf{p}) = \sum_{i=1}^N c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$.

2.2 Kinematics-Conditioned Latent Disentanglement

Existing frameworks tightly couple semantic features with geometric deformations, causing unnatural spatial warping during in-place semantic mutations (e.g., cauterization). We address this by disentangling semantic evolution from spatial displacement. Given a Gaussian at position $p \in \mathbb{R}^3$ and time t , a multi-resolution HexPlane encoder extracts a spatiotemporal feature $\mathbf{h}_{st} = \mathcal{H}(p, t)$. A shared base network processes $\mathbf{h}_{st}$ to yield a local dynamics representation $\mathbf{h}_{base}$. We project $\mathbf{h}_{base}$ into property-specific states to form a *latent kinematic descriptor* $\mathcal{Z}_{kin}$:

$$\mathcal{Z}_{kin} = \bigoplus_{k \in \mathcal{K}} \phi_k(\mathbf{h}_{base}) \in \mathbb{R}^{d_{kin}}, \quad \mathcal{K} = \{pos, scale, rot, opa\} \quad (1)$$

where $\phi_k(\cdot)$ denotes linear transformations. Rather than parameterizing the semantic feature $f_{sem}(t)$ via deformed geometry, two independent branches decode geometric deformations $\Delta \mathcal{G}_t$ and semantic updates $\Delta f_{sem,t}$ directly from $\mathcal{Z}_{kin}$:

$$\Delta \mathcal{G}_t = \Psi_{geo}(\mathcal{Z}_{kin}), \quad \Delta f_{sem,t} = \Psi_{sem}(\mathcal{Z}_{kin}) \quad (2)$$

This conditionally isolates geometric and semantic updates. During topology-altering events with static geometry but changing appearance, Ψ_{sem} freely triggers semantic shifts ($\Delta f_{sem} \neq \mathbf{0}$) without forcing false geometric displacements ($\Delta \mathcal{G} \rightarrow \mathbf{0}$).




Fig. 2. Overview of the proposed DeGenseGS framework. To prevent spurious geometric warping, a HexPlane encoder extracts shared kinematic latents to enable independent geometric and semantic decoding, thereby explicitly decoupling the optimization pathways. Subsequently, Angular-Aligned Semantic Distillation, together with a Rasterization-Native Semantic Extraction mechanism, strengthens feature grounding and facilitates accurate boundary recovery from rasterized feature maps.

2.3 Angular-Aligned Semantic Distillation

To distill 2D VLM knowledge, we optimize the angular similarity between rendered and ground-truth VLM features within their native hyperspherical space. Features $\mathbf{v}_{pred}(p)$ and $\mathbf{v}_{gt}(p)$ are L_2 -normalized to $\tilde{\mathbf{v}}$. The alignment loss $\mathcal{L}_{align}$ minimizes their angular discrepancy:

$$\mathcal{L}_{align} = \mathbb{E}_p [1 - \langle \tilde{\mathbf{v}}_{pred}(p), \tilde{\mathbf{v}}_{gt}(p) \rangle] \quad (3)$$

To enforce semantic consistency within anatomical structures, we add a region smoothness regularization $\mathcal{L}_{smooth} = \mathbb{E}_{R \in \Omega} [\frac{1}{|R|} \sum_{p \in R} \|\mathbf{v}_{pred}(p) - \bar{\mathbf{v}}_R\|_1]$, where R is a superpixel region and $\bar{\mathbf{v}}_R$ is its mean feature.

2.4 Rasterization-Native Semantic Extraction

To robustly extract semantics and boundaries, we process features directly in the 2D rendered space rather than via 3D primitives. First, the rasterizer projects 3D semantic features into a 2D feature map $\mathcal{F}_{sem}$, alongside the RGB image $\mathbf{I}_{rgb}$. To resolve visual-language ambiguity [3] in dense point-wise matching, we introduce **Region-aware Similarity Aggregation**.

An unsupervised segmentation algorithm decomposes the rendered image into superpixel regions $\Omega = \{R_1, \dots, R_K\}$. We aggregate the pixel-wise cosine

similarity between $\mathcal{F}_{sem}$ and a target text embedding $\mathbf{v}_{text}$ via mean pooling within each R_k :

$$S(R_k, \mathcal{T}) = \frac{1}{|R_k|} \sum_{p \in R_k} \left\langle \frac{\mathcal{F}_{sem}(p)}{\|\mathcal{F}_{sem}(p)\|_2}, \frac{\mathbf{v}_{text}}{\|\mathbf{v}_{text}\|_2} \right\rangle \quad (4)$$

This suppresses spurious background activations and resolves local ambiguities. Assigning this aggregated score to all pixels in R_k yields a coarse probability map $\mathbf{M}_{coarse}$. Finally, an edge-preserving Guided Filter G uses $\mathbf{I}_{rgb}$ as structural guidance to recover precise pixel-level boundaries: $\mathbf{M}_{final} = G(\mathbf{M}_{coarse}, \mathbf{I}_{rgb})$.

3 Experiments

3.1 Datasets and Evaluation Metrics

To evaluate DeGenseGS in complex surgical environments featuring severe non-rigid deformations and topology-altering events such as electrocautery, we conduct experiments on five sequences from the CholecSeg8k [5] benchmark and two sequences from the EndoVis18 [1] benchmark. The mean Intersection over Union metric is adopted as the primary evaluation criterion for 3D semantic segmentation accuracy.

3.2 Implementation Details

We conduct all experiments on a single NVIDIA RTX 3090 GPU. The model is optimized using the Adam optimizer with an initial learning rate of 1.6×10^{-3} . We employ a coarse-to-fine schedule, freezing the deformation parameters for the first 3,000 iterations. Standard Gaussian densification and pruning are performed every 100 iterations. In terms of efficiency, DeGenseGS renders at ~ 70 FPS, and a full text-prompted query with VLM matching and RNSE takes ~ 0.56 s per frame, on par with SurgTPGS [6] (~ 67 FPS and ~ 0.6 s per frame).

3.3 Quantitative and Qualitative Evaluation

We compare DeGenseGS against state-of-the-art GS frameworks and vision-language integrated methods, including LangSplat [18], OpenGaussian [23], DGD [11], FE-4DGS [13], and SurgTPGS [6]. To ensure a comprehensive assessment, variants equipped with different vision-language models such as SurgVLP [27] and CAT-Seg [4] are also evaluated.

As shown in Table 1 and Table 2, our method consistently outperforms all competitors. On CholecSeg8k, DeGenseGS establishes a new state-of-the-art score of 68.20%, achieving a substantial 14.74% absolute improvement over the best-performing baseline SurgTPGS at 53.46%. Similarly, on EndoVis18, we achieve 44.63%, significantly surpassing SurgTPGS at 39.09%. In addition, averaged over all evaluated scenes from both datasets, DeGenseGS also improves

Table 1. Quantitative results evaluated by mean Intersection over Union on the CholecSeg8k benchmark. **Bold**: best in column; underline: second best. Abbreviations include Abd. Wall for Abdominal Wall, Gallbl. for Gallbladder, L-hook for L-hook Electrocautery, Avg. for Average score, LS for LangSplat, and OG for OpenGaussian.

Methods	01_00080		01_00240		01_15019			12_15750				17_01803				Avg.
	Liver	Liver	Grasper	Abd. Wall	Grasper	Liver	Fat	Gallbl.	Grasper	L-hook	Liver	Abd. Wall	Fat	Grasper	Liver	
LangSplat	64.87	57.91	4.50	23.84	4.13	24.94	25.63	3.42	0	5.97	25.03	43.38	12.31	4.13	10.43	20.35
LS-SurgVLP	65.03	57.14	4.31	32.79	5.61	27.45	4.10	3.34	0	5.97	22.59	42.97	10.30	4.53	10.34	19.50
LS-CAT-Seg	77.51	73.69	4.96	91.05	58.29	<u>39.87</u>	75.83	5.93	0	23.81	30.60	72.93	24.00	14.46	12.84	39.37
OpenGaussian	2.64	0.08	0	0	0	0.42	0	0.68	0	8.05	0.21	0	1.63	0	0.19	0.93
OG-SurgVLP	1.76	0	0	0	0	0.57	0.19	0.60	0	0	0.15	13.02	1.42	2.24	0.70	1.38
OG-CAT-Seg	2.21	1.90	12.79	14.89	12.82	1.55	1.26	1.02	0	7.49	1.28	0	1.20	0	1.23	3.80
DGD	13.98	42.43	3.83	9.68	2.90	10.97	14.79	2.25	0.73	0	8.95	4.09	8.16	2.67	3.81	8.80
FE-4DGS	0.08	47.90	0.93	4.94	0.03	0.25	0.29	0.21	0	0	12.40	2.05	0.41	0.15	4.65	4.95
SurgTPGS	<u>88.70</u>	<u>81.27</u>	67.61	<u>97.18</u>	68.57	37.42	<u>73.98</u>	<u>7.63</u>	<u>16.34</u>	<u>54.04</u>	<u>34.12</u>	<u>89.26</u>	<u>28.24</u>	<u>58.74</u>	<u>15.11</u>	<u>53.46</u>
Ours	89.03	86.00	<u>61.05</u>	97.66	<u>62.59</u>	92.35	65.91	52.14	27.68	65.06	77.53	96.11	46.31	73.96	47.02	68.20

Table 2. Quantitative results evaluated by mean Intersection over Union on the EndoVis18 benchmark. **Bold**: best in column; underline: second best. Abbreviations include Avg. for Average score, LS for LangSplat, and OG for OpenGaussian.

Methods	Seq_5		Seq_9			Avg.
	Inst-Wrist	Kidney-Parenchyma	Inst-Shaft	Inst-Wrist	Inst-Clasper	
LangSplat	5.17	59.16	16.47	7.89	6.15	18.96
LS-SurgVLP	7.36	52.94	9.87	11.72	7.02	17.78
LS-CAT-Seg	17.16	54.42	38.49	<u>21.74</u>	0.11	26.38
OpenGaussian	0	1.62	0	0	0	0.32
OG-SurgVLP	1.13	0.54	0	0	0	0.33
OG-CAT-Seg	0	0.06	<u>29.38</u>	12.82	14.89	11.43
DGD	1.14	21.68	8.22	5.04	7.37	8.69
FE-4DGS	1.46	52.14	6.18	4.30	6.08	14.03
SurgTPGS	<u>43.29</u>	<u>71.98</u>	22.64	17.07	<u>40.48</u>	<u>39.09</u>
Ours	52.96	82.22	20.26	22.24	45.46	44.63

the underlying RGB reconstruction quality over SurgTPGS by +1.06 in PSNR, +0.030 in SSIM, and -0.0010 in LPIPS; the segmentation gain (+14.74% mIoU on CholecSeg8k) is substantially larger than the reconstruction gain, indicating that the improvement primarily originates from the proposed geometry-semantic decoupling rather than from better RGB reconstruction alone.

This performance leap is particularly prominent in structures subject to severe topological and appearance alterations. For instance, gallbladder segmentation accuracy in sequence 12_15750 surges from 7.63% to 52.14%, and instrument-wrist in Seq_5 improves from 43.29% to 52.96%. These gains empirically validate that disentangling semantic evolution from geometric warping effectively prevents feature collapse, preserving high spatial coherence amidst surgical dynamics. Qualitatively, as illustrated in Fig. 3, DeGenseGS demonstrates significantly enhanced geometric connectivity and precise semantic boundaries.




Fig. 3. Qualitative result on CholecSeg8k and EndoVis18 datasets.

Table 3. Ablation study of the proposed components. We report the mean Intersection over Union (mIoU %) across the overall CholecSeg8k and EndoVis18 datasets, along with a specific subset (Seq 12_15750) for fine-grained comparison.

S-G Decoupling	AASD	CholecSeg8k		EndoVis18	
		Seq 12_15750	Overall Average mIoU	Overall Average mIoU	
×	×	32.65	53.46	39.09	
×	✓	34.12	55.12	40.21	
✓	×	46.21	66.35	43.45	
✓	✓	48.57	68.20	44.63	

3.4 Ablation Study

To validate our contributions, Table 3 presents an ablation study evaluating our core components. Notably, our **Semantic-Geometric (S-G) Decoupling** encompasses both the Kinematics-Conditioned Latent Disentanglement (Sec. 2.2) and the Rasterization-Native Semantic Extraction (Sec. 2.4), which jointly decouple semantics from spatial warping at both the 4D feature level and the 2D rasterization level. For the rigidly coupled baseline (× for S-G Decoupling), we adopt the SurgTPGS architecture [6].

As shown, while the Angular-Aligned Semantic Distillation (AASD) objective alone provides a marginal performance improvement by regularizing the VLM latent space, its overall impact remains constrained by the baseline’s erroneous geometric warpings. In contrast, the complete S-G Decoupling framework acts as the primary driver of the substantial performance leap (e.g., boosting the CholecSeg8k average mIoU from 53.46% to 66.35%). This proves its crucial role

in mitigating feature collapse and gradient-induced geometric artifacts during complex topology-altering events.

4 Conclusion

We present DeGenseGS, the first text-promptable 4D Gaussian Splatting framework capable of separately characterizing spatial deformation and semantic evolution for fine-grained surgical interaction. By implementing a dual-branch architecture for structural decoupling and rasterization-native semantic extraction, our method addresses the fundamental challenge of gradient interference between geometry and semantics. This decoupled paradigm prevents spurious geometric warping while ensuring spatially coherent surgical understanding. The superior performance validates DeGenseGS as a robust foundation for reliable perception in autonomous robotic-assisted surgery.

Acknowledgments. This work was supported by Ministry of Education Tier 2 grant, Singapore (T2EP20224-0028), and Ministry of Education Tier 1 grant, Singapore (23-0651-P0001). This work was also supported by the China Scholarship Council under Grant No. 202506090084. Yimo Wang conducted this research as a CSC visiting PhD student at the National University of Singapore.

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

References

1. Allan, M., Kondo, S., Bodenstedt, S., Leger, S., Kadkhodamohammadi, R., Lungeno, I., et al.: 2018 Robotic Scene Segmentation Challenge. arXiv preprint arXiv:2001.11190 (2020)
2. Chen, J., Zhang, X., Hoque, M.I., Vasconcelos, F., Stoyanov, D., Elson, D.S., Huang, B.: SurgicalGS: Dynamic 3D Gaussian Splatting for Accurate Robotic-Assisted Surgical Scene Reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 572–582. Springer Nature Switzerland (2025)
3. Chen, Q., Yang, L., Chen, Y., Zhao, N., Lai, J., Shao, J., Xie, X.: Training-Free Class Purification for Open-Vocabulary Semantic Segmentation. In: IEEE/CVF International Conference on Computer Vision. pp. 23124–23134. IEEE (2025)
4. Cho, S., Shin, H., Hong, S., Arnab, A., Seo, P.H., Kim, S.: CAT-Seg: Cost Aggregation for Open-Vocabulary Semantic Segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4113–4123 (2024)
5. Hong, W.Y., Kao, C.L., Kuo, Y.H., Wang, J.R., Chang, W.L., Shih, C.S.: Cholec-Seg8k: A Semantic Segmentation Dataset for Laparoscopic Cholecystectomy Based on Cholec80. arXiv preprint arXiv:2012.12453 (2020)
6. Huang, Y., Bai, L., Cui, B., Yuan, K., Wang, G., Hoque, M.I., et al.: SurgTPGS: Semantic 3D Surgical Scene Understanding with Text Promptable Gaussian Splatting. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 584–594. Springer (2025)

7. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics* **42**(4), 139–1 (2023)
8. Kerr, J., Kim, C.M., Goldberg, K., Kanazawa, A., Tancik, M.: LERF: Language Embedded Radiance Fields. In: *IEEE/CVF International Conference on Computer Vision*. pp. 19729–19739. IEEE (2023)
9. Khan, U., Nawaz, U., Qayyum, A., Ashraf, S., Xie, Y., Khan, M.H., Qadir, J.: Surgical Scene Understanding in the Era of Foundation AI Models: A Comprehensive Review. *arXiv preprint arXiv:2502.14886* (2025)
10. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., et al.: Segment Anything. In: *IEEE/CVF International Conference on Computer Vision*. pp. 4015–4026. IEEE (2023)
11. Labe, I., Issachar, N., Lang, I., Benaïm, S.: DGD: Dynamic 3D Gaussians Distillation. In: *European Conference on Computer Vision*. pp. 361–378. Springer (2024)
12. Li, J., Skinner, G., Yang, G., Quaranto, B.R., Schwaitzberg, S.D., Kim, P.C., Xiong, J.: LLaVA-Surg: Towards Multimodal Surgical Assistant via Structured Surgical Video Learning. *arXiv preprint arXiv:2408.07981* (2024)
13. Li, K., Wang, J., Han, W., Zhao, D.: FeatureEndo-4DGS: Real-Time Deformable Surgical Scene Reconstruction and Segmentation with 4D Gaussian Splatting. In: *Machine Learning for Health Symposium*. pp. 1218–1234. PMLR (2026)
14. Li, Z., Chen, Z., Li, Z., Xu, Y.: Spacetime Gaussian Feature Splatting for Real-Time Dynamic View Synthesis. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8508–8520. IEEE (2024)
15. Liu, M., Han, Y., Wang, J., Wang, C., Wang, Y., Meijering, E.: LSKANet: Long Strip Kernel Attention Network for Robotic Surgical Scene Segmentation. *IEEE Transactions on Medical Imaging* **43**(4), 1308–1322 (2023)
16. Liu, Y., Li, C., Yang, C., Yuan, Y.: EndoGaussian: Real-Time Gaussian Splatting for Dynamic Endoscopic Scene Reconstruction. *arXiv preprint arXiv:2401.12561* (2024)
17. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment Anything in Medical Images. *Nature Communications* **15**(1), 654 (2024)
18. Qin, M., Li, W., Zhou, J., Wang, H., Pfister, H.: LangSplat: 3D Language Gaussian Splatting. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20051–20060. IEEE (2024)
19. Shan, J., Cai, Z., Hsieh, C.T., Han, L., Cheng, S.S., Wang, H.: Deformable Gaussian Splatting for Efficient and High-Fidelity Reconstruction of Surgical Scenes. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 10545–10551. IEEE (2025)
20. Wang, G., Bai, L., Nah, W.J., Wang, J., Zhang, Z., Chen, Z., et al.: Surgical-LVLM: Learning to Adapt Large Vision-Language Model for Grounded Visual Question Answering in Robotic Surgery. In: *ICLR 2025 Workshop on Foundation Models in the Wild* (2025)
21. Wang, Y., Long, Y., Fan, S.H., Dou, Q.: Neural Rendering for Stereo 3D Reconstruction of Deformable Tissues in Robotic Surgery. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 431–441. Springer (2022)
22. Wu, T., Miao, Y., Guo, J., Chen, Z., Zhao, S., Li, Z., et al.: EndoWave: Rational-Wavelet 4D Gaussian Splatting for Endoscopic Reconstruction. *arXiv preprint arXiv:2510.23087* (2025)

23. Wu, Y., Meng, J., Li, H., Wu, C., Shi, Y., Cheng, X., et al.: OpenGaussian: Towards Point-Level 3D Gaussian-Based Open Vocabulary Understanding. In: *Advances in Neural Information Processing Systems*. vol. 37, pp. 19114–19138 (2024)
24. Yang, C., Wang, K., Wang, Y., Yang, X., Shen, W.: Neural LerPlane Representations for Fast 4D Reconstruction of Deformable Tissues. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 46–56. Springer (2023)
25. Yang, S., Li, Q., Shen, D., Gong, B., Dou, Q., Jin, Y.: Deform3DGS: Flexible Deformation for Fast Surgical Scene Reconstruction with Gaussian Splatting. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 132–142. Springer (2024)
26. Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K., Finn, C.: Gradient Surgery for Multi-Task Learning. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 5824–5836 (2020)
27. Yuan, K., Srivastav, V., Yu, T., Lavanchy, J.L., Marescaux, J., Mascagni, P., Padoy, N.: Learning Multi-Modal Representations by Watching Hundreds of Surgical Video Lectures. *Medical Image Analysis* **105**, 103644 (2025)
28. Yue, W., Zhang, J., Hu, K., Xia, Y., Luo, J., Wang, Z.: SurgicalSAM: Efficient Class Promptable Surgical Instrument Segmentation. In: *AAAI Conference on Artificial Intelligence*. vol. 38, pp. 6890–6898. AAAI (2024)
29. Zhang, Y., Shen, Z., Jiao, R.: Segment Anything Model for Medical Image Segmentation: Current Applications and Future Directions. *Computers in Biology and Medicine* **171**, 108238 (2024)
30. Zhao, C., Huang, X., Yang, K., Wang, X., Wang, Q.: Generalizable 3D Gaussian Splatting for Novel View Synthesis. *Pattern Recognition* **161**, 111271 (2025)