

Endo-4DGX: Robust Endoscopic Scene Reconstruction and Illumination Correction with Gaussian Splatting

Yiming Huang^{1,2} ^{*}, Long Bai^{1,2} ^{*}, Beilei Cui^{1,2} ^{*}, Yanheng Li³ , Tong Chen⁴ , Jie Wang² , Jinlin Wu⁵ , Zhen Lei⁵ , Hongbin Liu⁵ , and Hongliang Ren^{1,2} ^{**}

¹ The Chinese University of Hong Kong, Hong Kong SAR, China

² Shenzhen Research Institute, CUHK, Shenzhen, China

³ City University of Hong Kong, Hong Kong, China

⁴ The University of Sydney, Sydney, NSW, Australia

⁵ Centre for Artificial Intelligence and Robotics, HKISI-CAS, Hong Kong, China
{yhuangdl, b.long, beileicui}@link.cuhk.edu.hk, hren@cuhk.edu.hk

Abstract. Accurate reconstruction of soft tissue is crucial for advancing automation in image-guided robotic surgery. The recent 3D Gaussian Splatting (3DGS) techniques and their variants, 4DGS, achieve high-quality renderings of dynamic surgical scenes in real-time. However, 3D-GS-based methods still struggle in scenarios with varying illumination, such as low light and over-exposure. Training 3D-GS in such extreme light conditions leads to severe optimization problems and devastating rendering quality. To address these challenges, we present Endo-4DGX, a novel reconstruction method with illumination-adaptive Gaussian Splatting designed specifically for endoscopic scenes with uneven lighting. By incorporating illumination embeddings, our method effectively models view-dependent brightness variations. We introduce a region-aware enhancement module to model the sub-area lightness at the Gaussian level and a spatial-aware adjustment module to learn the view-consistent brightness adjustment. With the illumination adaptive design, Endo-4DGX achieves superior rendering performance under both low-light and over-exposure conditions while maintaining geometric accuracy. Additionally, we employ an exposure control loss to restore the appearance from adverse exposure to the normal level for illumination-adaptive optimization. Experimental results demonstrate that Endo-4DGX significantly outperforms combinations of state-of-the-art reconstruction and restoration methods in challenging lighting environments, underscoring its potential to advance robot-assisted surgical applications. Our code is available at <https://github.com/lastbasket/Endo-4DGX>.

Keywords: 3D Reconstruction · Illumination Correction · Robotic Surgery.

^{*} Co-first authors.

^{**} Corresponding author.




Fig. 1. Our method achieves state-of-the-art performance for robust reconstruction and illumination correction in the surgical scene with varying illumination.

1 Introduction

Endoscopic procedures have emerged as a cornerstone of medical practice, particularly in the context of minimally invasive surgeries. In the recent development of robot-assisted minimally invasive surgeries, accurate reconstruction of the endoscopic scene is of the utmost significance. The novel reconstruction and rendering techniques like NeRF [17, 21, 28] and 3D Gaussian Splatting (3DGS) [12, 26] achieves superior performance for real-time renderings, aligning well with the surgical need for deformable surgical scenes reconstruction [10, 11, 15, 25].

Despite the notable achievements of existing methods, they face a formidable challenge in varying illumination within the surgical scenes. The endoscopic workspace is dynamic and unpredictable; illumination changes quickly due to factors such as the hardware issue of the light source and reflections from surgical instruments and tissue. Low-light conditions [16, 18] can make endoscope difficult to capture clear details of the soft tissue, while over-exposure [23] leads to information loss. Previous works [2, 5, 16, 19] have made progress in addressing 2D-level endoscopic illumination challenges. Garcia et al. [7] introduced a structure-aware network on endoscopic exposure correction (EC). Recently, EndoUIC [1] was proposed as a novel adaptive diffusion model for endoscopic EC.

However, due to the rapid variation of the surgical environment, 2D restoration methods [2, 7] fail to produce consistent results in real-time, leading to corrupted results for both appearance and geometry. While there exist 3DGS-based illumination correction solutions for general scenes [13, 27, 30], they are not applicable for surgical scenarios due to the lack of ability to model the sub-area and spatial-level illumination changes for deformable tissue, resulting in optimization failure of 3DGS. Besides, most general scene methods rely on pre-trained models [20] trained on daily life data [13, 31]. In this case, using such a pre-trained model in surgical scenes shall disrupt the training of Gaussians due to the domain gap, leading to unsatisfactory outcomes when directly applying a generic scene approach to the surgical domain. Therefore, we propose Endo-4DGX, specifically designed for endoscopic reconstruction and illumination correction as shown in Fig. 1. Endo-4DGX effectively models the view-dependent brightness variations

by incorporating *illumination embeddings*. To restore the reconstruction from challenging illumination, we develop two modules: The *region-aware enhancement* module for Gaussian-level sub-area refinement. The *spatial-aware adjustment* module for view-consistent exposure compensation across the entire image. Furthermore, we utilize *global exposure control* loss to achieve illumination-adaptive optimization, enabling a more adaptable and practical reconstruction in real surgical settings. Specifically, our contributions are listed as follows:

- We present Endo-4DGX, a novel endoscopic reconstruction method with illumination adaptive Gaussian Splatting. Our method achieves illumination correction and reconstruction in challenging uneven illumination.
- We introduce a region-aware enhancement module to resolve uneven lighting problems at the Gaussian level. Our region-aware module decodes the view-specific embedding to model the sub-area lightness changes.
- We design a spatial-aware adjustment module for spatial-level lightness adjustment, which focuses on the illumination refinement of the whole image.
- Our exhausted experiments on three real surgical datasets demonstrate the robustness of Endo-4DGX for surgical scene reconstruction under challenging illumination, providing a powerful tool for robot-assisted surgery.

2 Methodology

The overview of our proposed Endo-4DGX is shown in Fig 2. We first preprocess the input to obtain depth and illumination prior to initialization. Then, we train the Gaussians with our proposed *region-aware enhancement* and *spatial-aware adjustment* modules. Additionally, we normalize and optimize the rendering result with *global exposure control*. We support the varying illumination rendering and the illumination correction with illumination embedding-based rendering.

2.1 Preliminaries

4D Gaussian Splatting (4DGS) [22, 26] is an advanced technique for dynamic scene reconstruction, especially in endoscopic surgical applications [11, 15, 25]. We select 4DGS [22] as our baseline, which is built on 3D Gaussian Splatting [12] and represents a scene using 3D differentiable Gaussians for rapid rendering. Given a set of images and the corresponding camera poses, the deformable surgical scene can be represented by the 4DGS [11, 15, 22] $\mathcal{G}' = \{\mu', r', s', o', c\}$, where μ', r', s', o' are the deformable mean, rotation, scale, and opacity, c is the color. With the center depth d_i [11], the 2D covariance [32] Σ' and $\alpha = o' \cdot e^{\frac{1}{2}(x-\mu')^T(\Sigma')^{-1}(x-\mu')}$, the pixel-level color and depth is rendered by the alpha-blending [12] as:

$$C(x) = \sum_i \alpha_i c_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad D(x) = \sum_i \alpha_i d_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (1)$$




Fig. 2. Illustration of our proposed framework. We first preprocess the inputs for initialization, and then we train the Gaussians with our illumination embedding, region-aware enhancement, and spatial-aware adjustment for sub-area and global restoration.

2.2 Proposed Methodology: Endo-4DGX

Illumination Embedding and Region-Aware Enhancement. We propose illumination embedding and region-aware enhancement to address the challenges posed by varying illumination conditions. Before training, we pre-process all images $\{I\}_N$ by a forward-reverse illumination estimation algorithm [29] $\mathcal{D}(\cdot)$. We utilize the pre-processed output $\mathbf{p} = \mathcal{D}(I)$ as the lightness prior and classify the input images into two levels (Bright, Dark) for more precise and separated restoration. We define the illumination condition (IC) of each training image as:

$$IC(I) = \begin{cases} \text{Bright, if } \text{mean}(I) > \text{mean}(\mathbf{p}) \\ \text{Dark, if } \text{mean}(I) \leq \text{mean}(\mathbf{p}) \end{cases} \quad (2)$$

To deal with rapid illumination changes, we initialize trainable illumination embeddings $\mathbf{e} \in \mathbf{R}^{N \times k}$ for each specific time of view, in which N is the number of training images, and k is the embedding dimension which is set as 32 in our cases. During the training, all embeddings are updated with Gaussians jointly.

Inspired by [13], we construct our region-aware enhancement as an affine transformation for Gaussian-level illumination enhancement. Given a view-specific color C , We input the view-specific illumination embedding into a lightweight concealing network f_{region} . However, using a single network for both the bright and dark image correction will lead to significant issues of training imbalance and convergence failure. To handle this, we proposed a unique approach involving two separate networks with different IC (Bright and Dark), where each network is formed by three layers of MLPs and sigmoid activations. By dividing the training images into separate ICs, we set the concealing network to bright or dark each time for enhancement. We split the output from the concealing

network into two 1-channel enhancement parameters as β, γ . For a Gaussian’s view-dependent color, the enhanced color is calculated as:

$$c_{tone} = \beta \cdot c + \gamma, \mid \beta, \gamma = f_{region}(c, \mathbf{e}) \quad (3)$$

Our illumination enhancement module models and fine-tunes the sub-area lightness of Gaussians based on the learned illumination embeddings and view-specific characteristics, achieving more accurate and realistic surgical scene reconstructions in the presence of varying illumination.

Spatial-Aware Illumination Adjustment. While region-aware enhancement is useful for addressing illumination irregularities of individual elements or small regions within the scene, it lacks spatial-level adjustment ability, which makes it insufficient to tackle the broader challenges of image-level spatial illumination. To complement this, we further propose a spatial-aware illumination adjustment method. Inspired by [8], we perform spatial adjustment by gathering the rendering output C_{tone} from the regional enhanced color and generate a quadratic curve map with our spatial adjustment network $f_{spatial}$:

$$\tilde{C}_{tone} = C_{tone} + \delta \cdot C_{tone}(1 - C_{tone}) \mid \delta = f_{spatial}(\mathbf{e}), \quad (4)$$

where $f_{spatial}$ is a lightweight network with three layers of MLPs and tanh activation. By using this adjustment, we effectively fine-tune the illumination at the spatial level, making the overall scene more visually consistent.

Illumination-Adaptive Optimization. We introduce an illumination-adaptive optimization strategy to maintain stable training under varying illumination conditions. We utilize the L1 loss $\mathcal{L}_{l1}$ and Total Variation (TV) loss $\mathcal{L}_{tv}$ for the rendered color $\tilde{C}_{tone}$ and depth D with illumination embedding. Following [11], we supervise the render depth with the ground truth depth $\hat{D}$ from stereo matching or pre-trained model [24]. The loss of color and depth is defined as:

$$\mathcal{L}_{color} = |\tilde{C}_{tone} - \hat{C}| \odot M + \lambda_{tv} L_{tv}(\tilde{C}_{tone}), \quad (5)$$

$$\mathcal{L}_{depth} = \lambda_{depth} \left| \frac{D}{D_{max}} - \frac{\hat{D}}{\hat{D}_{max}} \right| \odot M + \lambda_{tv} L_{tv}(D), \quad (6)$$

where $\lambda_{depth}, \lambda_{tv}$ are the weight for depth and tv loss, M is the tool mask, and $\hat{C}$ is the ground truth color. Inspired by [6, 8], we also propose a global exposure control loss for the rendered color C without the illumination embedding:

$$\mathcal{L}_{Exposure} = \left\| \text{avgpool1d} \left(\frac{1}{3} \sum_{j \in R, G, B} C_j \right) - \mathbf{E} \right\|^2, \quad (7)$$

where $\mathbf{E}$ is the exposure level (0.6 is adopted), avgpool1d is a 1-dimensional pooling. The global exposure control loss ensures that the reconstructed scene has a consistent and normal exposure despite the inputs being over- or under-exposure. The final loss is represented as:

$$\mathcal{L} = \mathcal{L}_{color} + \mathcal{L}_{depth} + \mathcal{L}_{Exposure} \quad (8)$$

Table 1. Quantitative results on EndoNeRF-EC Dataset [21] for illumination correction. The first, second, and third values are highlighted. * indicates that the method fails with NaN loss during reconstruction. Our method shows overall superior performance over state-of-the-art baseline methods.

Method	FPS↑	GPU↓	Pulling			Cutting		
			PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
EndoUIC [1]+Deform3DGS [25]	2.02	8 GB	27.15	0.892	0.180	22.74	0.842	0.231
EndoUIC [1]+Endo-4DGS [11]	2.0	12 GB	28.79	0.929	0.097	23.30	0.890	0.116
EndoUIC [1]+EndoGaussian [15]	2.03	12 GB	27.49	0.904	0.244	23.05	0.867	0.257
CSEC [14]+Deform3DGS [25]	1.29	3.8 GB	20.03	0.855	0.194	20.18	0.809	0.251
CSEC [14]+Endo-4DGS [11]	1.27	7.8 GB	19.89	0.876	0.140	20.14	0.838	0.186
CSEC [14]+EndoGaussian [11]	1.29	7.8 GB	18.63	0.838	0.283	19.41	0.800	0.339
DarkGS* [30]	524	2 GB	6.04	0.001	0.655	6.66	0.002	0.635
Gaussian-DK* [27]	41.0	2 GB	13.13	0.705	0.528	10.26	0.543	0.579
WildGaussians [13]	108	2 GB	15.28	0.724	0.431	17.71	0.752	0.376
Ours	61	6.5 GB	34.94	0.946	0.048	34.06	0.950	0.043

Table 2. Quantitative results on StereoMIS [9] and C3VD [3] datasets for uneven illumination scene reconstruction. The first and second are highlighted. Our method surpasses all state-of-the-art methods.

Method	EndoNeRF-EC				StereoMIS				C3VD			
	Pulling		Cutting		P1_1		P1_2		Cecum_t2_b		Sigmoid_t2_a	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Deform3DGS [25]	30.63	0.922	26.86	0.831	28.55	0.846	32.28	0.887	33.07	0.898	27.92	0.749
Endo-4DGS [11]	28.51	0.909	26.88	0.850	26.98	0.817	32.48	0.890	30.55	0.878	32.03	0.830
EndoGaussian [15]	12.93	0.327	12.49	0.289	27.61	0.841	32.20	0.896	5.10	0.001	5.99	0.001
Ours	39.95	0.966	39.64	0.968	28.73	0.853	32.70	0.899	33.55	0.904	33.13	0.838

3 Experiments

3.1 Dataset

We split the training and testing sets as in [21] and evaluate all baselines on an exposure correction (EC) dataset and two datasets with uneven illumination:

EndoNeRF [21] Exposure Correction (EndoNeRF-EC) dataset includes two public video clips for robotic prostatectomy. The dataset is uniformly sampled as under-exposure, normal-exposure, and over-exposure in every 7 steps. We collaborate with professional photographers to simulate camera aperture adjustments using the Adobe Camera Raw SDK. The exposure values (EVs) are adjusted to analog the aperture changes. RAW images are rendered by different EV levels of three levels: under-exposure, over-exposure, and normal. Within each, the values are randomized, simulating real-world exposure mistakes.

StereoMIS [9] dataset includes surgical video sequences for in-vivo porcine subjects. We select two video clips of frames from P1: 16400-16697 (P1_1) and 7379-7499 (P1_2), which contain the illumination variance due to the tool occlusion and reflection. We split the training and testing sets following [11]. The ground truth depth is obtained by the stereo matching algorithm in OpenCV [4].




Fig. 3. Qualitative Result on EndoNeRF-EC [21] Dataset. Results for DarkGS and Gaussian-DK are not presented due to the reconstruction failure. Our method provides the best reconstruction and illumination correction results for challenging illumination.

C3VD [3] dataset is a video data collection designed for cardiovascular reconstruction. We utilize the *cecum_t2_b* and *sigmoid_t2_a* clips, including uneven illumination caused by the light source direction and occlusions.

3.2 Implementation Details

We evaluate the rendering quality by comparing the Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS). We also report the Frame Per Second (FPS) and GPU memory usage to compare the computation. All methods are trained on a single RTX4090 GPU. We utilize the learning rate 1.6×10^{-3} following [15] with Adam optimizer. For the hyperparameters, we adopt $\lambda_{depth}, \lambda_{tv} = 0.01$. To evaluate the illumination restoration, we utilize the first frame in the training data with normal-level illumination as the input embedding. While evaluating the uneven illumination reconstruction, we adopt the protocol in [13].

3.3 Results

We compare our method with three state-of-the-art Endoscopic Gaussian Splatting baselines [11, 15, 25] combined with state-of-the-art 2D illumination correction methods [1, 14], and additionally three Gaussian Splatting-based methods

Table 3. Ablation experiments on EndoNeRF-EC dataset [21]. We evaluate the PSNR and SSIM of the varying illumination reconstruction and illumination correction by removing: (i) illumination embedding, (ii) region-aware enhancement, and (iii) spatial-aware adjustment. The **first** and **second** result is highlighted.

Illumination Embedding	Region-Aware Enhancement	Spatial-Aware Adjustment	Pulling		Cutting		Pulling-Correction		Cutting-Correction	
			PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$
$\times$	$\times$	$\times$	26.95	0.892	27.59	0.856	22.88	0.770	20.53	0.734
$\times$	$\checkmark$	$\times$	10.40	0.535	14.64	0.624	12.66	0.579	13.96	0.712
$\times$	$\times$	$\checkmark$	24.56	0.840	30.53	0.895	23.46	0.815	19.93	0.692
$\checkmark$	$\checkmark$	$\times$	39.64	0.965	37.33	0.957	34.79	0.945	32.66	0.933
$\checkmark$	$\times$	$\checkmark$	34.87	0.952	34.52	0.939	28.44	0.910	25.09	0.865
$\times$	$\checkmark$	$\checkmark$	14.36	0.627	31.06	0.925	21.94	0.804	24.26	0.740
$\checkmark$	$\checkmark$	$\checkmark$	39.95	0.966	39.64	0.968	34.94	0.946	34.06	0.950

for low-light enhancement [27, 30] and uneven illumination reconstruction [13]. For evaluation of [30], we use the default configuration for da Vinci Xi system. Since there is no paired data for StereoMIS [9] and C3VD [3], we evaluate the illumination correction on the EndoNeRF-EC [21] dataset and evaluate the varying illumination reconstruction on all datasets. As demonstrated in Table 1, Endo-4DGX achieves state-of-the-art performance for the restoration task under challenging illumination conditions. Our method significantly outperforms existing reconstruction and restoration baselines for the illumination correction task on the EndoNeRF-EC dataset. Our method achieves PSNR/SSIM/LPIPS scores of 34.94/0.946/0.048 for the Pulling case and 34.06/0.950/0.043 for the Cutting case, surpassing the best baseline (EndoUIC [1]+Endo-4DGS [11]) by over 6 dB in PSNR. Notably, while maintaining real-time rendering at 61 FPS, our approach reduces GPU memory consumption by 46% compared to the combination of 2D restoration and 4DGS variants. In Table 2, despite illumination restoration, we conduct uneven illumination reconstruction on StereoMIS and C3VD datasets, where Endo-4DGX also demonstrates superior generalization with the highest PSNR. Additionally, our method achieves illumination correction under scenes without normal-level illumination embeddings by global exposure control. The qualitative result is shown in Fig 3, demonstrating the improvement in terms of visual quality. For more visualization, please refer to the supplementary.

Ablation results in Table 3 validate the necessity of each proposed component of our model. The results without illumination embedding drop significantly, even worse than the result without all modules, indicating that the integration of illumination embedding with region-aware enhancement and spatial-aware adjustment is necessary. Our full model outperforms the others without the proposed modules, confirming the critical role of each component in maintaining robust surgical scene reconstruction.

4 Conclusions

We propose Endo-4DGX, an illumination-aware Gaussian Splatting framework for endoscopic 3D reconstruction under challenging lighting conditions. Our

method achieves state-of-the-art performance for both reconstruction and illumination correction of dynamic lighting scenarios by unifying region-aware illumination enhancement and spatial-aware illumination adjustment. Experiments demonstrate our method outperforms existing methods while maintaining real-time rendering and robust generalization to various surgical datasets. This work advances practical applications in surgical navigation and robotic interventions, promoting the precision and safety of surgical procedures. Future research will extend the framework to broader clinical lighting challenges.

Acknowledgements. This work was supported by Hong Kong RGC CRF C4026-21G, RIF R4020-22, GRF 14211420, 14216020 & 14203323).

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

References

1. Bai, L., Chen, T., Tan, Q., Nah, W.J., Li, Y., He, Z., Yuan, S., Chen, Z., Wu, J., Islam, M., et al.: Endouic: Promptable diffusion transformer for unified illumination correction in capsule endoscopy. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 296–306. Springer (2024)
2. Bai, L., Chen, T., Wu, Y., Wang, A., Islam, M., Ren, H.: Llcaps: Learning to illuminate low-light capsule endoscopy with curved wavelet attention and reverse diffusion. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 34–44. Springer (2023)
3. Bobrow, T.L., Golhar, M., Vijayan, R., Akshintala, V.S., Garcia, J.R., Durr, N.J.: Colonoscopy 3d video dataset with paired depth from 2d-3d registration. Medical Image Analysis p. 102956 (2023)
4. Bradski, G.: The opencv library. Dr. Dobb’s Journal of Software Tools (2000)
5. Chen, T., Lyu, Q., Bai, L., Guo, E., Gao, H., Yang, X., Ren, H., Zhou, L.: Lightdiff: Surgical endoscopic image low-light enhancement with t-diffusion. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 369–379. Springer (2024)
6. Cui, Z., Gu, L., Sun, X., Ma, X., Qiao, Y., Harada, T.: Aleth-nerf: Illumination adaptive nerf with concealing field assumption. In: Proceedings of the AAAI Conference on Artificial Intelligence (2024)
7. Garcia-Vega, A., Espinosa, R., Ramírez-Guzmán, L., Bazin, T., Falcón-Morales, L., Ochoa-Ruiz, G., Lamarque, D., Daul, C.: Multi-scale structural-aware exposure correction for endoscopic imaging. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2023)
8. Guo, C., Li, C., Guo, J., Loy, C.C., Hou, J., Kwong, S., Runmin, C.: Zero-reference deep curve estimation for low-light image enhancement. CVPR (2020)
9. Hayoz, M., Hahne, C., Gallardo, M., Candinas, D., Kurmann, T., Allan, M., Sznitman, R.: Learning how to robustly estimate camera pose in endoscopic videos. International Journal of Computer Assisted Radiology and Surgery pp. 1185–1192 (2023)
10. Huang, Y., Cui, B., Bai, L., Chen, Z., Wu, J., Li, Z., Liu, H., Ren, H.: Advancing dense endoscopic reconstruction with gaussian splatting-driven surface normal-aware tracking and mapping. arXiv preprint arXiv:2501.19319 (2025)

11. Huang, Y., Cui, B., Bai, L., Guo, Z., Xu, M., Islam, M., Ren, H.: Endo-4dgs: Endoscopic monocular scene reconstruction with 4d gaussian splatting. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 197–207. Springer (2024)
12. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (2023)
13. Kulhanek, J., Peng, S., Kukulova, Z., Pollefeys, M., Sattler, T.: WildGaussians: 3D gaussian splatting in the wild. *NeurIPS* (2024)
14. Li, Y., Xu, K., Hancke, G.P., Lau, R.W.: Color shift estimation-and-correction for image enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 25389–25398 (2024)
15. Liu, Y., Li, C., Yang, C., Yuan, Y.: Endogaussian: Gaussian splatting for deformable surgical scene reconstruction. *arXiv preprint arXiv:2401.12561* (2024)
16. Ma, Y., Liu, Y., Cheng, J., Zheng, Y., Ghahremani, M., Chen, H., Liu, J., Zhao, Y.: Cycle structure and illumination constrained gan for medical image enhancement. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23. pp. 667–677. Springer (2020)
17. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
18. Moghtaderi, S., Yaghoobian, O., Wahid, K.A., Lukong, K.E.: Endoscopic image enhancement: Wavelet transform and guided filter decomposition-based fusion approach. *Journal of Imaging* **10**(1), 28 (2024)
19. Mou, E.: Global and local enhancement of low-light endoscopic images-[v2] (2023)
20. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
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, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20310–20320 (June 2024)
23. Yang, K.F., Cheng, C., Zhao, S.X., Yan, H.M., Zhang, X.S., Li, Y.J.: Learning to adapt to light. *International Journal of Computer Vision* **131**(4), 1022–1041 (2023)
24. Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10371–10381 (2024)
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. Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20331–20341 (2024)

27. Ye, S., Dong, Z.H., Hu, Y., Wen, Y.H., Liu, Y.J.: Gaussian in the dark: Real-time view synthesis from inconsistent dark images using gaussian splatting. In: Computer Graphics Forum. vol. 43, p. e15213. Wiley Online Library (2024)
28. Zha, R., Cheng, X., Li, H., Harandi, M., Ge, Z.: Endosurf: Neural surface reconstruction of deformable tissues with stereo endoscope videos. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 13–23. Springer (2023)
29. Zhang, Q., Nie, Y., Zheng, W.S.: Dual illumination estimation for robust exposure correction. In: Computer graphics forum. vol. 38, pp. 243–252. Wiley Online Library (2019)
30. Zhang, T., Huang, K., Zhi, W., Johnson-Roberson, M.: Darkgs: Learning neural illumination and 3d gaussians relighting for robotic exploration in the dark. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 12864–12871. IEEE (2024)
31. Zuo, X., Samangouei, P., Zhou, Y., Di, Y., Li, M.: Fmgs: Foundation model embedded 3d gaussian splatting for holistic 3d scene understanding. International Journal of Computer Vision pp. 1–17 (2024)
32. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Surface splatting. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 371–378 (2001)