

BridgeSplat: Bidirectionally Coupled CT and Non-Rigid Gaussian Splatting for Deformable Intraoperative Surgical Navigation

Maximilian Fehrentz^{1,2}, Alexander Winkler², Thomas Heiliger²,
Nazim Haouchine³, Christian Heiliger², and Nassir Navab¹

¹ Computer Aided Medical Procedures, TU Munich, Munich, Germany

² Department of General, Visceral, and Transplantation Surgery, Hospital of the LMU Munich, Ludwig-Maximilians-Universität (LMU), Munich, Germany

³ Harvard Medical School, Brigham and Women's Hospital, Boston, MA, USA

Abstract. We introduce BridgeSplat, a novel approach for deformable surgical navigation that couples intraoperative 3D reconstruction with preoperative CT data to bridge the gap between surgical video and volumetric patient data. Our method rigs 3D Gaussians to a CT mesh, enabling joint optimization of Gaussian parameters and mesh deformation through photometric supervision. By parametrizing each Gaussian relative to its parent mesh triangle, we enforce alignment between Gaussians and mesh and obtain deformations that can be propagated back to update the CT. We demonstrate BridgeSplat's effectiveness on visceral pig surgeries and synthetic data of a human liver under simulation, showing sensible deformations of the preoperative CT on monocular RGB data. Code, data, and additional resources can be found at <https://maxfehrentz.github.io/ct-informed-splatting/>.

Keywords: Non-Rigid Gaussian Splatting · CT-Coupled Intraoperative Reconstruction · Deformable Surgical Navigation

1 Introduction

Laparoscopic surgery has shown the potential to lower complication rates, shorten hospital stays, and be cost-effective while maintaining R0 resection rates and showing no adverse effects on 90-day mortality [3]. However, it also comes with challenges, as tactile feedback is limited, spatial understanding is demanding, handling is challenging, and learning curves are flat. To mitigate those issues, Augmented Reality (AR) has been proposed to overlay relevant information from preoperative scans with little interruption to the clinical workflow. The potential for AR-guided (minimally invasive) interventions in visceral surgery has been demonstrated for the liver [1], kidney [10], pancreas [21], and spleen [18]. Although the results are promising, those methods either stop at a rigid registration or require interruption of the surgical workflow for surgeon interaction (e.g., outlining contours). Continuous non-rigid registration is still an open challenge.

Thinking from first principles, we would like a system that can reconstruct the intraoperative scene in real time while also estimating the camera pose. Therefore, much effort has been directed at non-rigid SLAM for surgery [14,12,17]. Despite all the progress on SLAM, as concluded in [13], tracking of the camera is still desirable, if not necessary, to move towards clinically feasible systems.

For posed cameras in surgical scenes, methods capable of 4D reconstruction for dynamic photorealistic novel view synthesis have emerged. Those methods are grounded in NeRF [20] and 4D Gaussian Splatting (4DGS) [24,9,8,22,7]. Prior work has been done based on both NeRF and 4DGS to go beyond mere reconstruction and towards intelligence, tracking the surgical scene [4,6]. However, NeRF does not have an explicit representation, and 4DGS is usually not bound to a surface. Attempting a monocular single-view reconstruction commonly leads to artifacts like floaters in the view frustum that allow overfitting on the individual observations but do not guarantee a meaningful 3D reconstruction of tissue over time, despite high image reconstruction quality.

In [19], the authors propose to extract a preoperative 3D model, equip it with a biomechanical model, and perform an initial rigid registration based on an intraoperative point cloud acquired with stereo depth. To allow for a non-rigid registration, they track the registered structure and deform the biomechanical model, using optical flow. Also building on biomechanical models, the authors in [2] propose a CNN for non-rigid registration.

Although some of the introduced methods leverage *selected* 3D information from preoperative scans, none of them is *directly coupled* to the preoperative scan, using the entire 3D prior for guiding the reconstruction and enabling direct deformation propagation to the CT.

Contribution We deform a preoperative CT from an intraoperative *monocular* RGB video. By rigging Gaussians to a registered CT mesh, we constrain the ill-posed monocular reconstruction problem through a canonical space prior, resolving the depth ambiguity. This reframes the problem of non-rigid registration from a non-rigid reconstruction to a non-rigid tracking problem with the preoperative CT as template. Unlike existing 4D Gaussian Splatting methods, our approach is bound to patient-specific anatomy, and its deformations are directly coupled to the CT.

2 Method

2.1 Problem Formulation

We start with a preoperative mesh \mathbf{M} of the abdominal cavity, obtained from the preoperative CT. At the beginning of surgery, a rigid registration \mathbf{T}_{CT} is performed to align the CT (and thus \mathbf{M}) with the patient. We then introduce a set of 3D Gaussians, which are rigged onto the registered mesh $\mathbf{T}_{CT}\mathbf{M}$. The 3D Gaussians have learnable parameters $\Theta_t = \{\mu_t, \alpha_t, \Sigma_t, \mathbf{c}_t\}$: means μ_t , transparencies α_t , covariance matrices Σ_t , and spherical harmonic coefficients




Fig. 1: We initialize 3D Gaussians by rigging them onto the registered mesh, extracted from the preoperative CT. Given the first camera pose, we identify the visible Gaussians, derive the visible vertices, and subsample them to initialize a sparse deformation field. For all consecutive frames, we optimize for the mesh vertices, the deformation vectors, and the appearance properties of the 3D Gaussians. Deformations are propagated to the CT.

\mathbf{c}_t . Additionally, a deformation field θ_t is defined to capture mesh deformations over time.

During surgery, a tracked laparoscope provides images \mathbf{I}_t from known camera poses \mathbf{P}_t . Our main objective is to minimize a photometric energy term $\mathbf{E}_{\text{photo}}$ between the observed image \mathbf{I}_t and a rendered image $\hat{\mathbf{I}}_t$ via 3D Gaussian Splatting of the *deformed* mesh and its rigged 3D Gaussians. Formally, we seek to solve:

$$\arg \min_{\theta_t, \Theta_t} \mathbf{E}_{\text{photo}} = \arg \min_{\theta_t, \Theta_t} \left\| \mathbf{I}_t - \hat{\mathbf{I}}_t(\mathbf{P}_t, \theta_t(\mathbf{T}_{CT}\mathbf{M}), \Theta_t) \right\|. \quad (1)$$

where θ_t deforms the registered mesh $\mathbf{T}_{CT}\mathbf{M}$. Solving this over time yields dynamic 3D Gaussians, commonly referred to as 4D Gaussian Splatting.

2.2 Coupling 3D Gaussians and Mesh

Unlike conventional 3D Gaussian Splatting methods that treat the means μ_t and covariances Σ_t as unconstrained parameters, we tightly anchor our Gaussians to the mesh, inspired by GaussianAvatars [11]. Specifically, each triangle f^i of the deformed mesh $\theta_t(\mathbf{T}_{CT}\mathbf{M})$ holds one or more Gaussians whose means are expressed via optimizable barycentric coordinates:

$$\mu_t^i = \sum_{k=1}^3 \mathbf{b}_t^{i,k} \mathbf{v}_t^{i,k}, \quad \text{with} \quad \sum_{k=1}^3 \mathbf{b}_t^{i,k} = 1, \quad \mathbf{b}_t^{i,k} \geq 0, \quad (2)$$

where $\mathbf{v}_t^{i,k} = \theta_t(\mathbf{v}_{\text{init}}^{i,k})$ are the vertices of triangle f^i after deformation, and $\mathbf{b}_t^i = (b_t^{i,1}, b_t^{i,2}, b_t^{i,3})$ are the barycentric coordinates for the i -th Gaussian. By

enforcing nonnegative $\mathbf{b}_t^{i,k}$ that sum to 1, the mean $\boldsymbol{\mu}_t^i$ cannot leave the parent face. This is in contrast to GaussianAvatars, where Gaussians are allowed to wander off the surface for enhanced visual fidelity. In our case, however, we try to estimate mesh deformations from single-view monocular RGB input. Therefore, we have to ensure that what we render is closely related to what the mesh actually represents, otherwise, Gaussians can satisfy the photometric energy by roaming and stretching freely beyond their parent triangles without the mesh having to deform.

For $\boldsymbol{\Sigma}_t^i$, we decompose the covariance into a rotation (orientation) and an anisotropic scale. In particular, we fix the Gaussian’s orientation \mathbf{q}_t^i to align with the triangle’s normal, thus keeping the Gaussian aligned with the surface of its parent triangle. The corresponding scale vector \mathbf{s}_t^i has only *two* freely optimizable components, the component in normal direction is fixed to a small constant. We further cap the remaining scales to keep them within a factor of the initial scale, which is computed from a mesh-dependent heuristic. This prevents the Gaussians from stretching significantly beyond their parent triangles.

2.3 Deformation Field

As shown in Equation (2), the means of the Gaussians are parametrized w.r.t to \mathbf{M} by the barycentric coordinates and a sparse deformation field θ_t for each timestep t that is acting on the mesh. We mostly follow the deformation field parametrization as proposed in [6]. However, instead of defining control points based on anchor 3D Gaussians, we choose control vertices. For the first frame, we perform a single render to check for visible 3D Gaussians. We can trace those back to their parent faces and obtain all visible vertices. We then choose a random subset of anchor vertices $\mathbf{a}_1, \dots, \mathbf{a}_k \in \mathbf{M}$, where each anchor vertex has a learnable deformation $\delta_1, \dots, \delta_k \in \mathbb{R}^3$. It provides continuous deformations through interpolation as described in [6]. Note that choosing initially visible vertices as anchors comes with the downside of modeling deformations only in a fixed field of view and its adjacency. For our short sequences, this is not an issue. For more extensive camera sweeps, the method can be easily extended to continuously check for visibility and add control vertices as the camera moves.

2.4 Non-rigid Regularization

To regularize the deformations, we employ an As-Rigid-As-Possible (ARAP) term \mathbf{E}_{ARAP} [15]. Note that this would be infeasible in a mesh-free approach, since ARAP inherently uses mesh topology and cannot be applied directly to unconstrained Gaussian Splatting. Following [6], we additionally penalize relative position changes of neighboring vertices and deformations in currently invisible areas.

For clarity, we omit the time dependency in the notation below. ARAP operates on a set of deformed vertices $\{\mathbf{v}'_1, \dots, \mathbf{v}'_n\} \subset \mathbb{R}^3$ of a deformed mesh $\mathbf{M}_{\text{deform}}$. For each vertex $\mathbf{v}'_i \in \mathbb{R}^3$, ARAP measures the discrepancy between the edges in $\mathbf{M}_{\text{deform}}$ and the corresponding edges in the original mesh \mathbf{M} under an

estimated rotation $\mathbf{R}_i \in \text{SO}(3)$ for each vertex. This encourages local transformations to remain as rigid as possible. In its standard formulation, ARAP is usually driven by a small number of known control points, and the minimization of \mathbf{E}_{ARAP} proceeds via a flip-flop scheme: first, one estimates the per-vertex rotations $\{\mathbf{R}_1, \dots, \mathbf{R}_n\} \subset \text{SO}(3)$ using SVD, then one updates the vertex positions $\{\mathbf{v}'_1, \dots, \mathbf{v}'_n\} \subset \mathbb{R}^3$ by directly solving a linear system. In our approach, we first set $\{\mathbf{v}'_1, \dots, \mathbf{v}'_n\} = \theta(\mathbf{M})$. Then, we simply compute the rotations $\{\mathbf{R}_1, \dots, \mathbf{R}_n\}$ via SVD as proposed in the original ARAP. However, we do not solve for the optimal vertex positions directly. Instead, we compute $\nabla_{\mathbf{v}'_i} \mathbf{E}_{\text{ARAP}}$ as derived in the original ARAP paper and use this term in our gradient-based optimization.

3 Experiments and Results

We evaluate our method on two datasets. Given that intraoperative data does not come with ground truth deformations, we simulate tool-tissue interaction sequences to evaluate our method quantitatively. To demonstrate its ability to operate on clinical data and deform a CT, we use a second dataset from visceral pig surgeries. All experimental protocols were approved by the local Ethical Committee on Animal Experimentation. Since no comparable methods are available, we cannot compare against a state-of-the-art. Whereas other 4DGS methods are primarily concerned with high-quality novel view synthesis, we focus on the deformation of the mesh and CT and therefore do not compare image reconstruction metrics. Results are best viewed in video format, we refer the reader to the supplementary material.

Quantitative Evaluation on Simulated Data We generate five synthetic tool-tissue interactions on a human liver, using DejaVu simulation [5]. DejaVu provides realistic surgical scenes by transferring the appearance of real surgical images to physics-based simulations. The deformations are computed using the Finite Element Method in SOFA [16]. Examples are shown in Figure 2.




Fig. 2: Example sequence of a simulated DejaVu [5] scene: on top, the textured mesh, tool, and a realistic background, underneath the 3D model colored based on stress.

For each sequence, we perform 45 simulation steps. We have an initial mesh and at each step in the simulation, we extract the image without the surgical tool, corresponding camera parameters, and the underlying mesh ground truth. We generate five different force profiles, yielding five simulated sequences with 45 images, camera parameters, and mesh ground truths each. To cover a wide range of interactions, we choose a different predominant manipulation in each sequence: pushing into the liver (SimIn), pulling to the left (SimLeft), pulling down (SimDown), dragging in a circular motion (SimCircular), and pulling up (SimUp).

We run our method on the five sequences and in each simulation step, we compute the Euclidean distance between each vertex in our deformed mesh and the corresponding vertex in the ground truth mesh. For each simulated sequence, we compute the average Euclidean distance, its standard deviation, and the maximum Euclidean distance. As reported in Table 1, we stay within a 5 mm error margin and observe only slight deviations in accuracy across the different interactions, indicating that our method can handle in-plane (e.g., SimLeft) and out-of-plane deformations (e.g., SimIn). A selection of tracked deformations on the DeJaVu dataset is shown in Figure 3.

Table 1: Quantitative evaluation on the DeJaVu dataset with Euclidean distance and maximum error (in mm) on five simulated tool-tissue interaction sequences.

Metric	SimIn	SimDown	SimLeft	SimCircular	SimUp
Euclidean Distance (mm)	0.11 ± 0.19	0.14 ± 0.27	0.09 ± 0.16	0.10 ± 0.19	0.12 ± 0.24
Max Euclidean Distance (mm)	3.17	4.29	2.78	4.15	2.84




Fig. 3: Exemplary results on a simulated DeJaVu [5] sequence. Rows top to bottom: input image with surgical tool and background image removed, rendered image from Gaussian Splatting, and deformed mesh. The first column shows the first step of the sequence, where some of the Gaussians still have random colors. See video in the supplementary material.

Qualitative Evaluation on Clinical Data To demonstrate the effectiveness of our method on clinical data, we use a dataset of two visceral pig surgeries with a preoperative post-insufflation CT, initial rigid registration, a tracked laparoscope, and corresponding 2D monocular RGB videos. Surgical tools were masked out using SurgicalSAM [23]. The first sequence from the first surgery shows a tool-tissue interaction, whereas the second one shows a breathing motion without external manipulation. Because there is no deformed ground truth, as measuring intraoperative deformation would require intraoperative scanning, we demonstrate qualitative results.

Figure 4 shows results on the tool-tissue interaction. The surgical tool presses down on the stomach, leading to the deformations tracked in the deformed mesh and CT (bottom rows). Although the difference in pressure applied between B and C appears marginal, we successfully track the increasing deformation, the difference being visible in both mesh and CT. For D and E, the tool is repositioned, now pulling the stomach back. Note how this backward pull is visible in the respective deformed CTs by a deformation to the *right*, since the laparoscope captures the scene roughly in the coronal plane, whereas the CT slice is shown in the sagittal plane.




Fig. 4: Examples from a tool-tissue interaction sequence. Rows from top to bottom: input images acquired by the laparoscope, images rendered with our method, deforming mesh, and deforming CT. Arrows in A point to the deforming anatomical structure, circles in B highlight the area of deformation, and arrows in E visualize the pull direction. See video in the supplementary material.

Since we rely on an initial registration, we are also subject to registration errors. There is a slight misalignment between the structure being manipulated in the input images (stomach) and the deforming mesh. Also note that we initialize our mesh from a post-insufflation CT, ensuring that the initial mesh and the intraoperative state early in surgery are still very similar. *Initializing* the method from video sequences later in the procedure cannot be handled by the method if significant deformation has already occurred or large structures have already disappeared entirely. Regarding the deformation of the CT, we can only capture surface deformations by deforming our mesh. A more realistic deformation of deeper layers would require biomechanical modeling. Apart from those caveats, the method does not work in real time yet. Depending on image resolution, processing of a frame takes several seconds on an NVIDIA RTX A4000. Given the strides being made in the 3D computer vision community around 4DGS, we are confident that this problem will be solved soon. Despite those limitations, the deformations of the stomach in the CT are sensible and correspond to the tool-tissue interaction as confirmed by a board-certified surgeon.

Coming from our clinical motivation, we also demonstrate potential dynamic overlays in Figure 5 for the sequence of the second surgery, displaying breathing motion. The overlays improve over the initial rigid registration, demonstrating the method’s ability to pick up on subtle deformations.




Fig. 5: Top row: overlays based on rigid registration. Bottom row: deformed overlays with our method. Deformations are subtle (see video).

4 Conclusion

In this paper, we have demonstrated for the first time that a direct coupling between CT and 4D Gaussian Splatting is possible to deform preoperative volumetric patient data during surgery. By blending the concepts of 4D Gaussian Splatting and deforming 3D meshes, we have shown that a mesh of the abdominal cavity can serve as an effective intermediary to capture changes observed

through the laparoscope and relay them directly to the CT. We have evaluated our method on synthetic and clinical data, demonstrating that its accuracy meets clinical needs on simulated data and performs well under real conditions in visceral pig surgeries. Strikingly, this is possible on 2D RGB data even without the need for stereo laparoscopes or any depth information.

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

References

- Bertrand, L.R., Abdallah, M., Espinel, Y., Calvet, L., Pereira, B., Ozgur, E., Pezet, D., Buc, E., Bartoli, A.: A case series study of augmented reality in laparoscopic liver resection with a deformable preoperative model. *Surgical Endoscopy* **34**(12), 5642–5648 (Dec 2020). <https://doi.org/10.1007/s00464-020-07815-x>
- Docea, R., Pfeiffer, M., Müller, J., Krug, K., Hardner, M., Riedel, P., Menzel, M., Kolbinger, F.R., Frohneberg, L., Weitz, J., et al.: A laparoscopic liver navigation pipeline with minimal setup requirements. In: 2022 IEEE Biomedical Circuits and Systems Conference (BioCAS). pp. 578–582. IEEE (2022)
- Fretland, A., Dagenborg, V.J., Bjørnelv, G.M.W., Kazaryan, A.M., Kristiansen, R., Fagerland, M.W., Hausken, J., Tønnessen, T.I., Abildgaard, A., Barkhatov, L., Yaqub, S., Røsok, B.I., Bjørnbeth, B.A., Andersen, M.H., Flatmark, K., Aas, E., Edwin, B.: Laparoscopic Versus Open Resection for Colorectal Liver Metastases: The OSLO-COMET Randomized Controlled Trial. *Annals of Surgery* **267**(2), 199 (Feb 2018). <https://doi.org/10.1097/SLA.0000000000002353>
- Gerats, B.G., Wolterink, J.M., Mol, S.P., Broeders, I.A.: Neural fields for 3d tracking of anatomy and surgical instruments in monocular laparoscopic video clips. *Healthcare Technology Letters* **11**(6), 411–417 (2024)
- Haouchine, N., Stoyanov, D., Roy, F., Cotin, S.: Dejavu: Intra-operative simulation for surgical gesture rehearsal. In: Medical Image Computing and Computer-Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part II 20. pp. 523–531. Springer (2017)
- Hayoz, M., Hahne, C., Kurmann, T., Allan, M., Beldi, G., Candinas, D., Márquez-Neila, P., Sznitman, R.: Online 3D reconstruction and dense tracking in endoscopic videos . In: proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. vol. LNCS 15006. Springer Nature Switzerland (October 2024)
- 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: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, vol. 15006, pp. 197–207. Springer Nature Switzerland, Cham (2024). https://doi.org/10.1007/978-3-031-72089-5_19, series Title: Lecture Notes in Computer Science
- Li, C., Feng, B.Y., Liu, Y., Liu, H., Wang, C., Yu, W., Yuan, Y.: EndoSparse: Real-Time Sparse View Synthesis of Endoscopic Scenes using Gaussian Splatting (Jul 2024), <http://arxiv.org/abs/2407.01029>, arXiv:2407.01029
- Liu, H., Liu, Y., Li, C., Li, W., Yuan, Y.: LGS: A Light-weight 4D Gaussian Splatting for Efficient Surgical Scene Reconstruction (Jun 2024), <http://arxiv.org/abs/2406.16073>, arXiv:2406.16073

10. Piana, A., Amparore, D., Sica, M., Volpi, G., Checcucci, E., Piramide, F., De Cillis, S., Busacca, G., Scarpelli, G., Sidoti, F.: Automatic 3D augmented-reality robot-assisted partial nephrectomy using machine learning: our pioneer experience. *Cancers* **16**(5), 1047 (2024), publisher: MDPI
11. Qian, S., Kirschstein, T., Schoneveld, L., Davoli, D., Giebenhain, S., Nießner, M.: GaussianAvatars: Photorealistic Head Avatars with Rigged 3D Gaussians. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20299–20309. IEEE, Seattle, WA, USA (Jun 2024). <https://doi.org/10.1109/CVPR52733.2024.01919>
12. Rodriguez, J.J.G., Montiel, J.M.M., Tardos, J.D.: NR-SLAM: Non-Rigid Monocular SLAM (Aug 2023), <http://arxiv.org/abs/2308.04036>, arXiv:2308.04036
13. Shu, H., Liu, M., Seenivasan, L., Gu, S., Ku, P.C., Knopf, J., Taylor, R., Unberath, M.: Seamless Augmented Reality Integration in Arthroscopy: A Pipeline for Articular Reconstruction and Guidance (Oct 2024), <http://arxiv.org/abs/2410.00386>, arXiv:2410.00386
14. Song, J., Wang, J., Zhao, L., Huang, S., Dissanayake, G.: MIS-SLAM: Real-time Large Scale Dense Deformable SLAM System in Minimal Invasive Surgery Based on Heterogeneous Computing (Mar 2018), <http://arxiv.org/abs/1803.02009>, arXiv:1803.02009
15. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: Symposium on Geometry processing. vol. 4, pp. 109–116. Citeseer (2007)
16. Talbot, H., Haouchine, N., Peterlik, I., Dequidt, J., Duriez, C., Delingette, H., Cotin, S.: Surgery Training, Planning and Guidance Using the SOFA Framework. In: Hege, H.C., Ropinski, T. (eds.) Eurographics 2015 - Dirk Bartz Prize. The Eurographics Association (2015). <https://doi.org/10.2312/egm.20151028>
17. Tang, X., Tao, H., Qian, Y., Yang, J., Feng, Z., Wang, Q.: Real-time deformable SLAM with geometrically adapted template for dynamic monocular laparoscopic scenes. *International Journal of Computer Assisted Radiology and Surgery* **19**(7), 1375–1383 (Jul 2024). <https://doi.org/10.1007/s11548-024-03174-3>
18. Tao, H.S., Lin, J.Y., Luo, W., Chen, R., Zhu, W., Fang, C.H., Yang, J.: Application of Real-Time Augmented Reality Laparoscopic Navigation in Splenectomy for Massive Splenomegaly. *World Journal of Surgery* **45**(7), 2108–2115 (Jul 2021). <https://doi.org/10.1007/s00268-021-06082-8>
19. Wang, E., Liu, Y., Tu, P., Taylor, Z.A., Chen, X.: Video-based Soft Tissue Deformation Tracking for Laparoscopic Augmented Reality-based Navigation in Kidney Surgery. *IEEE Transactions on Medical Imaging* pp. 1–1 (2024). <https://doi.org/10.1109/TMI.2024.3413537>, conference Name: IEEE Transactions on Medical Imaging
20. 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)
21. Wu, X., Wang, D., Xiang, N., Pan, M., Jia, F., Yang, J., Fang, C.: Augmented reality-assisted navigation system contributes to better intraoperative and short-time outcomes of laparoscopic pancreaticoduodenectomy: a retrospective cohort study. *International Journal of Surgery* **109**(9), 2598 (Sep 2023). <https://doi.org/10.1097/JS9.0000000000000536>
22. Yang, S., Li, Q., Shen, D., Gong, B., Dou, Q., Jin, Y.: Deform3DGS: Flexible Deformation for Fast Surgical Scene Reconstruction with Gaussian Splatting (May 2024), <http://arxiv.org/abs/2405.17835>, arXiv:2405.17835

23. Yue, W., Zhang, J., Hu, K., Xia, Y., Luo, J., Wang, Z.: Surgicalsam: Efficient class promptable surgical instrument segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 6890–6898 (2024)
24. Zhu, L., Wang, Z., Cui, J., Jin, Z., Lin, G., Yu, L.: EndoGS: Deformable Endoscopic Tissues Reconstruction with Gaussian Splatting (Jul 2024), <http://arxiv.org/abs/2401.11535>, arXiv:2401.11535