



Endo-FASt3r: Endoscopic Foundation model Adaptation for Structure from motion

Mona Sheikh Zeinoddin^{1,2}, Mobarak I. Hoque^{1,4}, Zafer Tandogdu^{3,6}, Greg L. Shaw³, Matthew J. Clarkson^{1,4}, Evangelos B. Mazomenos^{1,4}, and Danail Stoyanov^{1,5}

¹ UCL Hawkes Institute, University College London, UK
{mona.zeinoddin.22; e.mazomenos}@ucl.ac.uk

² Institute of Health Informatics, University College London, UK

³ Dept of Urology, University College London Hospitals, UK

⁴ Dept of Medical Physics & Biomedical Engineering, University College London, UK

⁵ Dept of Computer Science, University College London, UK

⁶ Division of Surgery and Interventional Science, University College London, UK

Abstract. Accurate depth and camera pose estimation is essential for achieving high-quality 3D visualisations in robotic-assisted surgery. Despite recent advancements in foundation model adaptation to monocular depth estimation of endoscopic scenes via self-supervised learning (SSL), no prior work has explored their use for pose estimation. These methods rely on low rank-based adaptation approaches, which constrain model updates to a low-rank space. We propose Endo-FASt3r, the first monocular SSL depth and pose estimation framework that uses foundation models for both tasks. We extend the Reloc3r relative pose estimation foundation model by designing Reloc3rX, introducing modifications necessary for convergence in SSL. We also present DoMoRA, a novel adaptation technique that enables higher-rank updates and faster convergence. Experiments on the SCARED dataset show that Endo-FASt3r achieves a substantial 10% improvement in pose estimation and a 2% improvement in depth estimation over prior work. Similar performance gains on the Hamlyn and StereoMIS datasets reinforce the generalisability of Endo-FASt3r across different datasets. Our code is available at: https://github.com/Mona-ShZeinoddin/Endo_FAST3r.git.

Keywords: Foundation model Adaptation · Depth and Pose estimation.

1 Introduction

Accurate depth and camera pose estimation are prerequisites to gaining a deeper geometric understanding of robotic-assisted surgery (RAS) scenes [18]. However, the unique properties of RAS such as varying illumination, textureless areas, and frequent occlusions complicate these tasks [4]. Although deep learning-based methods are a common approach [1], their development is limited by the lack of large-scale annotated endoscopic datasets, driving the adoption of self-supervised learning (SSL). With the rise of foundation models [27,16], recent works have

explored their potential in RAS [5,29]. However, directly applying them to RAS is suboptimal [29], inspiring the introduction of adaptation techniques. The potential of these techniques and the need for SSL have motivated the adaptation of foundation models to RAS using SSL. While many SSL monocular depth and pose estimation methods in RAS have utilised foundation models in their depth module [5,29], none have explored their use in the pose module. Leveraging a foundation model in the pose module can address key challenges, such as variations in camera movement and occlusions, which often lead to suboptimal performance in current CNN-based pose modules that focus mainly on local patterns [21,30]. Recent works on parameter-efficient fine-tuning techniques (PEFT) in SSL monocular depth estimation of RAS scenes [29,5] have been limited by the use of low-rank updates driven by the low-rank adaptation technique (LoRA) [10]. While LoRA’s approach of decomposing the weight update matrix into low-rank matrices is effective, it constrains adaptability, particularly in cases with significant domain shifts, such as RAS, where higher-rank updates are essential for optimal performance [11]. We introduce Endo-FASt3r, a novel SSL monocular depth and pose estimation framework that relies on foundation models for both tasks, making Endo-FASt3r the first to adapt a foundation model for pose estimation in RAS scenes. For the pose module, we design Reloc3rX, extending the Reloc3r [7] foundation model for robust convergence in SSL. We also present DoMoRA, a novel adaptation technique that advances recent PEFTs [11,13] by incorporating both low-rank and full-rank updates while benefiting from faster convergence. Our main contributions and findings include:

1. We design Reloc3rX by extending [7], which addresses the scale sensitivity issue of monocular SSL-based depth and pose estimation. Our work is the first to use a foundation model for pose estimation in RAS.
2. We propose DoMoRA, a novel adaptation technique that enhances previous PEFTs by addressing the necessity of both full-rank and low-rank updates while improving convergence through weight decomposition.
3. Experimental results on three public datasets, SCARED, Hamlyn, and, StereoMIS show a substantial improvement ranging from 7%–10% in the absolute trajectory error metric for pose estimation and a 2% improvement in the absolute relative error metric for depth estimation over state-of-the-art (SOTA) methods, highlighting the benefits of Endo-FASt3r.

2 Methods

2.1 Preliminaries

Depth Anything (DA) V2: Previously trained on 62M urban scene images, DA V2 employs a 12-layer DINO V2-based encoder [16] and a decoder with convolutional head and neck modules [27]. This work aims to better adapt DA V2 for depth estimation in RAS through the proposed adaptation technique, DoMoRA.

Reloc3r: To directly regress the relative pose between image pairs [7], Reloc3r

processes each image I_i through 24 ViT encoder blocks, producing features F_1 and F_2 for each image. The decoder includes 12 ViT blocks, each integrating a cross-attention layer between its self-attention and feed-forward layers to enhance spatial alignment between F_1 and F_2 . The two branches are fully symmetrical with shared weights. The decoder produces features G_1 and G_2 , which are then fed into a pose regression head comprising 2 feed-forward layers to regress the 9D intermediate representation of rotation and 3D translation from I_1 to I_2 and I_2 to I_1 :

$$\hat{R}_{I_1, I_2}^{(3 \times 3)}, t_{I_1, I_2}^{(3 \times 1)} = \text{Head}(G_1), \quad \hat{R}_{I_2, I_1}^{(3 \times 3)}, t_{I_2, I_1}^{(3 \times 1)} = \text{Head}(G_2) \quad (1)$$

The 9D intermediate rotation representation is then orthogonalised using singular value decomposition into a 3×3 rotation matrix, denoted as function *Orthogonal*.

$$R_{I_1, I_2}^{(3 \times 3)} = \text{Orthogonal}(\hat{R}_{I_1, I_2}^{(3 \times 3)}), \quad R_{I_2, I_1}^{(3 \times 3)} = \text{Orthogonal}(\hat{R}_{I_2, I_1}^{(3 \times 3)}) \quad (2)$$

Reloc3r was trained on 6M images in a supervised manner via the angular distance between the predicted and ground truth poses as the supervision signal. However, we have empirically verified that directly using Reloc3r in an SSL framework for RAS scenes causes divergence. Our work focuses on redesigning Reloc3r’s pose head to account for scale sensitivity, making Endo-FAST3r the first framework to adapt a foundation model for pose estimation in RAS scenes.

MoRA: Improving upon LoRA-based techniques by using a square matrix as opposed to low-rank matrices, MoRA enhances the capacity of the weight update matrix [11]. Given the pre-trained weight $W_0 \in \mathbb{R}^{d \times k}$, MoRA allocates a square matrix $M \in \mathbb{R}^{r \times r}$ with rank r and defines the new layer with MoRA as:

$$h(x) = W_0 x + f_{\text{decomp}} \left(M f_{\text{comp}}(x) \right) \quad (3)$$

In which $f_{\text{comp}} : \mathbb{R}^k \rightarrow \mathbb{R}^r$ and $f_{\text{decomp}} : \mathbb{R}^r \rightarrow \mathbb{R}^d$ are non-parameterized functions that compress the input dimension from k to r and decompress the output dimension from r to d respectively.

DoRA: Proposed to improve LoRA by drawing influence from Weight Normalization [20], DoRA decomposes the weight matrix into its magnitude and direction to accelerate optimisation [13]. The weight decomposition of $W \in \mathbb{R}^{d \times k}$ can be written as $W = m \frac{V}{\|V\|_c}$ where $m \in \mathbb{R}^{1 \times k}$ is the magnitude vector, $V \in \mathbb{R}^{d \times k}$ is the directional matrix, with $\|\cdot\|_c$ being the column-wise vector norm of a matrix. At initialisation, given the pre-trained weight W_0 , we set $m = \|W_0\|_c$ and $V = W_0$. LoRA with rank r is then applied only to the directional component while the magnitude component is kept trainable. The updated weight W' is:

$$W' = m \frac{V + \Delta V}{\|V + \Delta V\|_c} = m \frac{W_0 + \underline{BA}}{\|W_0 + \underline{BA}\|_c} \quad (4)$$

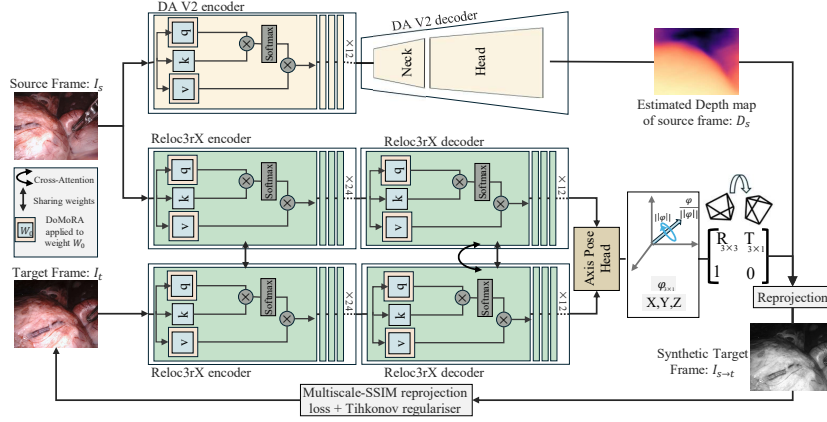


Fig. 1: EndoFast3r SSL monocular depth and pose estimation pipeline based on the two foundation models, DA V2 and Reloc3rX.

where ΔV is the incremental directional update learned by multiplying two low-rank matrices B and A , with the underlined parameters being trainable. The new layer with DoRA adaptation $h(x)$ can be defined as $h(x) = W'x$.

Self-Supervised Depth and Pose Estimation: The reprojection loss-based approach [8] is a widely used SSL method in monocular depth and pose estimation, especially in RAS [22]. It comprises a depth module, which estimates the depth map of the source frame I_s , and a pose module, which predicts the camera movement between the source frame I_s and target frame I_t . Given the depth prediction D_s for I_s , and R and t , the rotation and translation matrices describing the camera movement, the synthetic target frame can be derived using the reprojection function π as $I_{s \rightarrow t} = \pi(D_s, K, R, t, I_s)$, in which K is the camera intrinsics. A multi-scale SSIM (MS_SSIM)-based [25] reprojection loss proposed in [29] is used to measure the dissimilarity between the original and synthetic image:

$$L_{\text{reproj}} = \alpha \left(1 - \text{MS_SSIM}(I_t, I_{s \rightarrow t}) \right) + \beta |I_t - I_{s \rightarrow t}| \quad (5)$$

To address the varying illumination conditions and occlusions in RAS scenes, we have followed the approach in [22] which introduces the Tikhonov regulariser loss terms. The total loss L_{ssl} can be written as $L_{\text{ssl}} = L_{\text{reproj}} + L_{\text{tikhonov}}$.

2.2 Proposed Framework: Endo-FASt3r

As presented in Fig. 1, Endo-FASt3r structures an SSL framework for monocular endoscopic depth and pose estimation by incorporating DA V2 in the depth module and Reloc3rX in the pose module. We propose a novel adaptation method, DoMoRA, to better adapt both DA V2 and Reloc3rX to RAS scenes.

DoMoRA: To benefit from the faster convergence of DoRA and the full-rank update space of MoRA, we introduce DoMoRA. Following Eqs. 4, 3 and given the pre-trained weight $W_0^{d \times k}$, rank r , magnitude m , DoRA low-rank matrices $B^{d \times r}, A^{r \times k}$ and MoRA square matrix $M^{r \times r}$, the new layer with DoMoRA is:

$$h(x) = \left(\underline{m} \frac{W_0 + \underline{BA}}{\|W_0 + \underline{BA}\|_c} \right) x + f_{\text{decomp}} \left(M f_{\text{comp}}(x) \right) \quad (6)$$

In which f_{comp} and f_{decomp} follow the rotation and truncation methods[11]. Endo-FAST3r incorporates DoMoRA in the query and value matrices of all transformer blocks (Fig. 2). Inspired by the hierarchical feature learning nature of neural networks, we adopt the vector rank approach presented in [29].




Fig. 2: Illustration of the DoMoRA Transformer block, in which matrices M, m, A, B are trainable while pre-trained q, k, v matrices are frozen. In addition, DA V2 decoder neck, head and Reloc3rX pose head are also trainable.

Reloc3rX: We design Reloc3rX, an extension of Reloc3r tailored for RAS. Unlike supervised methods, SSL-based monocular depth and pose estimation suffers from scale ambiguity [24]. This is caused by the scale-sensitive Tikhonov regulariser [22] loss term, which leads to training instability and even divergence. We empirically verified that directly applying Reloc3r, pre-trained in a supervised manner on urban scenes, results in divergence during the early stages of training in RAS, due to scale misalignment. To address this, the pose regression head in Eq. 1 was modified to predict rotation in the 3D axis angle space instead of the 9D rotation matrix space, eliminating the need for orthogonalisation in Eq. 2. To ensure alignment with RAS scale, particularly at the early stages of training, the predicted axis-angle and translation vectors are scaled by a factor of 0.001. After scaling, the resulting axis angle $\phi_{I_1, I_2}^{(3 \times 1)}$ is converted to a rotation matrix via the Rodrigues rotation formula [15].

$$\phi_{I_1, I_2}^{(3 \times 1)}, t_{I_1, I_2}^{(3 \times 1)} = \text{Head}(G_1^{(T \times d)}) \quad (7)$$

Unlike Reloc3r, Reloc3rX computes the pose only from the source image to the target image, as the reverse is unnecessary for our SSL framework.

Training and Evaluation framework: In the training phase, given the source frame I_s and the target frame I_t , the depth module operates on the source frame I_s and produces its corresponding depth map. Meanwhile, the pose module will operate on both frames to estimate their relative camera pose. The synthetic target frame is then generated via reprojection. The Tikhonov regulariser and reprojection loss are used for optimisation. In the evaluation phase, each module operates independently to produce the desired depth map and camera pose.

3 Implementation details

Training: Endo-FAST3r utilises the Adam optimiser with a learning rate of $1e-4$ decaying every 10 steps. Training was performed on an A100 GPU for 10 epochs with a batch size of 4 for 20 hours.

Datasets and Evaluation Protocol: Endo-FAST3r has been trained and evaluated on the SCARED dataset [2], captured with a da Vinci Xi from porcine cadavers using structured light for depth collection. Following the split in [22], 15,351 frames were used for training, 1,705 for validation, and 551 for depth evaluation. Meanwhile pose evaluation utilised two trajectories of length 410 and 833 frames defined in [22]. To examine generalisability, we used the Hamlyn dataset following [19] for the depth module. For the pose module, we extracted frames 9,780–10,980 from sequence P3 of the StereoMIS dataset [9], which exhibit significant tissue deformation and camera motion in an *in vivo* setup. StereoMIS ground-truth poses were derived from the kinematics of the endoscope. The Depth evaluation metrics, mainly the absolute relative error (AbsRel), were adopted from [22]. Before evaluation, depth maps were scaled following [31]. The Absolute trajectory error (ATE) [14] was used for pose evaluation.

4 Results

We primarily compare Endo-FAST3r against SOTA reprojection-based SSL approaches for monocular depth and pose estimation [3,8,17,28,29] in RAS. To assess the generalisability of non-endoscopic foundation models DA V2 and Reloc3r, we also include their zero-shot performance. We use the reported metrics for all comparison methods in Table 1 [29,6]. To ensure fairness, we excluded SSL methods with extra modules beyond depth and pose in their reprojection, such as camera intrinsics estimation or image intrinsic decomposition, which add extra self-supervision signals [6,12]. Even if included, Endo-FAST3r still excels significantly in pose estimation and is comparable or superior in depth estimation, proving its high performance without requiring additional parameters.

Table 1: Quantitive evaluation on the SCARED and Hamlyn datasets. ATE-T1/T2 denote ATE for trjactorries 1/2. Total/Train denote the total/trainable (millions) depth module parameters. Speed (ms) is reported for the depth module.

| | Method | AbsRel↓ | SqRel↓ | RMSE↓ | $\delta \uparrow$ | ATE-T1 ↓ | ATE-T2 ↓ | Total Train | Speed |
|--------|--------------------------|--------------|--------------|---------------|-------------------|---------------|---------------|-------------|-------|
| SCARED | DeFeat-Net [23] | 0.077 | 0.792 | 6.688 | 0.941 | 0.1765 | 0.0995 | 14.8 14.8 | - |
| | SC-SfMLearner [3] | 0.068 | 0.645 | 5.988 | 0.957 | 0.0767 | 0.0509 | 14.8 14.8 | - |
| | Monodepth2 [8] | 0.069 | 0.577 | 5.546 | 0.948 | 0.0769 | 0.0554 | 14.8 14.8 | - |
| | Endo-SfM [17] | 0.062 | 0.606 | 5.726 | 0.957 | 0.0759 | 0.0500 | 14.8 14.8 | - |
| | AF-SfMLearner [22] | 0.059 | 0.435 | 4.925 | 0.974 | 0.0757 | 0.0501 | 14.8 14.8 | 8.0 |
| | Yang et al. [28] | 0.062 | 0.558 | 5.585 | 0.962 | 0.0723 | 0.0474 | 2.0 2.0 | - |
| | Zero-Shot DA V2 [27] | 0.091 | 1.056 | 7.601 | 0.916 | - | - | - | - |
| | Zero-Shot Reloc3r [27] | - | - | - | - | 0.0938 | 0.0735 | - | - |
| | DARES [29] | 0.052 | 0.356 | 4.483 | 0.980 | 0.0752 | 0.0498 | 24.9 2.88 | 15.6 |
| | EndoFAST3r (Ours) | 0.051 | 0.354 | 4.480 | 0.998 | 0.0702 | 0.0438 | 24.9 2.93 | 19.1 |
| Hamlyn | Endo Depth & Motion [19] | 0.185 | 5.424 | 16.100 | 0.732 | - | - | - | - |
| | AF-SfMLearner [22] | 0.168 | 4.440 | 13.870 | 0.770 | - | - | 14.8 14.8 | 7.7 |
| | EndoFAST3r (Ours) | 0.166 | 4.529 | 13.718 | 0.778 | - | - | 24.9 2.93 | 19.1 |

Pose estimation: Evaluation on the rigid SCARED dataset shows Endo-FAST3r outperforming all SOTA methods, improving by 6.64% and 12.04% - average 9.34% - over the second-best approach, DARES (with a CNN-based pose module), for trajectories 1 and 2 respectively (Table 1). Fig. 3 presents this improvement on trajectory 2, where occlusion by surgical fluid is effectively addressed by Endo-FAST3r. To assess generalisability to a non-rigid case with RAS-like tissue deformation and camera motion, we evaluated on the StereoMIS dataset, comparing against DARES, using its public code (Fig. 5). Endo-FAST3r outperforms DARES by 7.13% (Table 3). The pose module operates in real time, with a speed of 26ms. Other works do not report this quantity.




Fig. 3: Qualitative pose estimation comparison on the SCARED dataset

Depth Estimation: Evaluation on the SCARED dataset shows Endo-FAST3r outperforming all SOTA methods, surpassing the second-best approach, DARES by 2%. While both use DA V2, DARES employs a LoRA-based adaptation, highlighting the advantages of DoMoRA over LoRA-based methods. Further testing on the Hamlyn dataset confirms Endo-FAST3r’s effectiveness, with a 1.19% improvement over the next-best approach (Table 1). Depth estimation outputs

from both datasets are shown in Fig. 4. Notably, Endo-FAST3r only has 2.93 million trainable parameters (11.7% of total parameters), and maintains a real-time inference speed of $19.1ms$, slightly increased compared to SOTA methods.




Fig. 4: Qualitative depth estimation comparison on the SCARED and Hamlyn datasets - Endo-Fast3r better captures edges demonstrating its robustness.

We also present 3D reconstructions on the SCARED dataset in Fig. 5, which shows Endo-FAST3r better capturing darker areas with fewer artifacts.




Fig. 5: Qualitative comparison with the second-best approach, DARES: SCARED 3D reconstruction (left) - StereoMIS pose estimation (right).

Ablation Studies: Table 2 shows that DoRA and Reloc3rX improve pose estimation over the LoRA-based DARES (in the first row), while maintaining the same depth estimation accuracy, and convergence time despite Endo-FAST3r’s increased number of parameters, thanks to DoRA’s faster convergence properties. Using MoRA’s full-rank matrices alone proves suboptimal. By integrating foundation models in both modules and leveraging both full-rank and low-rank updates, Endo-FAST3r enhances performance in both tasks while matching DARES in convergence time, despite having more parameters.

Table 2: Ablation studies on the SCARED dataset.

| DoRA | MoRA | DoMoRA | Reloc3rX | AbsRel↓ | SqRel↓ | RMSE↓ | $\delta \uparrow$ | ATE-T1↓ | ATE-T2↓ |
|------|------|--------|----------|--------------|--------------|--------------|-------------------|---------------|---------------|
| × | × | × | × | 0.052 | 0.366 | 4.522 | 0.979 | 0.0752 | 0.0498 |
| ✓ | × | × | ✓ | 0.052 | 0.387 | 4.475 | 0.986 | 0.0735 | 0.0486 |
| × | ✓ | × | ✓ | 0.057 | 0.465 | 5.195 | 0.971 | 0.0748 | 0.0505 |
| × | × | ✓ | ✓ | 0.051 | 0.354 | 4.480 | 0.998 | 0.0702 | 0.0438 |

Table 3: StereoMIS results.

| Method | ATE↓ |
|--------------------------|---------------|
| DARES [29] | 0.0715 |
| EndoFAST3r (Ours) | 0.0664 |

5 Conclusion

We present Endo-FASt3r, a self-supervised monocular depth and pose estimation framework for RAS. We introduce Reloc3rX, the first development of a foundation model for pose estimation in endoscopic scenes. We also introduce DoMoRA, a novel adaptation technique that enables both low-rank and full-rank updates while benefitting from faster convergence. As the first SSL-based method that uses foundation models for both depth and pose estimation, Endo-FASt3r achieves a substantial improvement ranging from 7% – 10% in the ATE metric for pose estimation and 2% in the AbsRel metric in depth estimation over prior methods, as shown across three public datasets. Future work will explore the use of video foundation models [26] to improve temporal stability.

Acknowledgments. This work was supported in whole, or in part, by the the UKRI Centre for Doctoral Training (CDT) in AI-enabled Healthcare Systems [EP/S201612/1], the Department of Science, Innovation and Technology (DSIT), the Royal Academy of Engineering under the Chair in Emerging Technologies scheme, and the Engineering and Physical Sciences Research Council (EPSRC) under grants [EP/W00805X/1, EP/Z534754/1, UKRI145]. For the purpose of open access, the author has applied a CC BY public copyright licence to any author accepted manuscript version arising from this submission.

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

References

1. Ali, S.: Where do we stand in ai for endoscopic image analysis? deciphering gaps and future directions. *NPJ Digit. Med.* **5**(1), 184 (2022)
2. Allan, M., Mcleod, J., Wang, C., Rosenthal, J.C., Hu, Z., Gard, N., Eisert, P., Fu, K.X., Zeffiro, T., Xia, W., et al.: Stereo correspondence and reconstruction of endoscopic data challenge. *arXiv* (2021)
3. Bian, J., Li, Z., Wang, N., Zhan, H., Shen, C., Cheng, M.M., Reid, I.: Unsupervised scale-consistent depth and ego-motion learning from monocular video. *Neural Inf. Process. Syst.* **32** (2019)
4. Chen, L., Tang, W., John, N.W., Wan, T.R., Zhang, J.J.: Slam-based dense surface reconstruction in monocular minimally invasive surgery and its application to augmented reality. *Comput. Methods Programs Biomed* **158**, 135–146 (2018)
5. Cui, B., Islam, M., Bai, L., Ren, H.: Surgical-dino: adapter learning of foundation models for depth estimation in endoscopic surgery. *Int. J. Comput. Assist. Radiol. Surg.* pp. 1–8 (2024)
6. Cui, B., Islam, M., Bai, L., Wang, A., Ren, H.: Endodac: Efficient adapting foundation model for self-supervised depth estimation from any endoscopic camera. In: *MICCAI*. pp. 208–218. Springer (2024)
7. Dong, S., Wang, S., Liu, S., Cai, L., Fan, Q., Kannala, J., Yang, Y.: Reloc3r: Large-scale training of relative camera pose regression for generalizable, fast, and accurate visual localization. *arXiv* (2024)

8. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: CVPR. pp. 3828–3838 (2019)
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* **18**(7), 1185–1192 (2023)
10. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. *arXiv* (2021)
11. Jiang, T., Huang, S., Luo, S., Zhang, Z., Huang, H., Wei, F., Deng, W., Sun, F., Zhang, Q., Wang, D., et al.: Mora: High-rank updating for parameter-efficient fine-tuning. *arXiv* (2024)
12. Li, B., Liu, B., Zhu, M., Luo, X., Zhou, F.: Image intrinsic-based unsupervised monocular depth estimation in endoscopy. *IEEE J. Biomed. Health Inform* (2024)
13. Liu, S.Y., Wang, C.Y., Yin, H., Molchanov, P., Wang, Y.C.F., Cheng, K.T., Chen, M.H.: Dora: Weight-decomposed low-rank adaptation. *arXiv* (2024)
14. Mur-Artal, R., Montiel, J.M.M., Tardos, J.D.: Orb-slam: a versatile and accurate monocular slam system. *IEEE Trans. Robot.* **31**(5), 1147–1163 (2015)
15. Murray, R.M., Li, Z., Sastry, S.S.: A mathematical introduction to robotic manipulation. CRC press (2017)
16. 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* (2023)
17. Ozyoruk, K.B., Gokceler, G.I., Bobrow, T.L., Coskun, G., Incetan, K., Almalioglu, Y., Mahmood, F., Curto, E., Perdigoto, L., Oliveira, M., et al.: Endoslam dataset and an unsupervised monocular visual odometry and depth estimation approach for endoscopic videos. *Med. Image Anal* **71**, 102058 (2021)
18. Qian, L., Wu, J.Y., DiMaio, S.P., Navab, N., Kazanzides, P.: A review of augmented reality in robotic-assisted surgery. *IEEE trans. med. robot. bionics* **2**(1), 1–16 (2019)
19. Recasens, D., Lamarca, J., Fácil, J.M., Montiel, J., Civera, J.: Endo-depth-and-motion: Reconstruction and tracking in endoscopic videos using depth networks and photometric constraints. *IEEE Robot. Autom. Lett.* **6**(4), 7225–7232 (2021)
20. Salimans, T., Kingma, D.P.: Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Neural Inf. Process. Syst.* **29** (2016)
21. Schmidt, A., Mohareri, O., DiMaio, S., Yip, M.C., Salcudean, S.E.: Tracking and mapping in medical computer vision: A review. *Med. Image Anal* p. 103131 (2024)
22. Shao, S., Pei, Z., Chen, W., Zhu, W., Wu, X., Sun, D., Zhang, B.: Self-supervised monocular depth and ego-motion estimation in endoscopy: Appearance flow to the rescue. *Med. Image Anal* **77**, 102338 (2022)
23. Spencer, J., Bowden, R., Hadfield, S.: Defeat-net: General monocular depth via simultaneous unsupervised representation learning. In: CVPR. pp. 14402–14413 (2020)
24. Wang, C., Buenaposada, J.M., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: CVPR. pp. 2022–2030 (2018)
25. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: Asilomar Conf. Signals Syst. Comput. vol. 2, pp. 1398–1402. Ieee (2003)
26. Yang, H., Huang, D., Yin, W., Shen, C., Liu, H., He, X., Lin, B., Ouyang, W., He, T.: Depth any video with scalable synthetic data. *arXiv* (2024)
27. Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. *arXiv* (2024)

28. Yang, Z., Pan, J., Dai, J., Sun, Z., Xiao, Y.: Self-supervised lightweight depth estimation in endoscopy combining cnn and transformer. *IEEE Trans. Med. Imaging* (2024)
29. Zeinoddin, M.S., Lena, C., Qu, J., Carlini, L., Magro, M., Kim, S., De Momi, E., Bano, S., Grech-Sollars, M., Mazomenos, E., et al.: Dares: Depth anything in robotic endoscopic surgery with self-supervised vector-lora of the foundation model. In: *ECCV*. Springer (2024)
30. Zhong, J., Ren, H., Chen, Q., Zhang, H.: A review of deep learning-based localization, mapping and 3d reconstruction for endoscopy. *Journal of Micro and Bio Robotics* **21**(1), 1–25 (2025)
31. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: *CVPR*. pp. 1851–1858 (2017)