




Surgical-DINO: adapter learning of foundation models for depth estimation in endoscopic surgery

Beilei Cui¹ · Mobarakol Islam² · Long Bai¹ · Hongliang Ren^{1,3} 

Received: 12 January 2024 / Accepted: 16 February 2024
© The Author(s) 2024

Abstract

Purpose Depth estimation in robotic surgery is vital in 3D reconstruction, surgical navigation and augmented reality visualization. Although the foundation model exhibits outstanding performance in many vision tasks, including depth estimation (e.g., DINOv2), recent works observed its limitations in medical and surgical domain-specific applications. This work presents a low-ranked adaptation (LoRA) of the foundation model for surgical depth estimation.

Methods We design a foundation model-based depth estimation method, referred to as Surgical-DINO, a low-rank adaptation of the DINOv2 for depth estimation in endoscopic surgery. We build LoRA layers and integrate them into DINO to adapt with surgery-specific domain knowledge instead of conventional fine-tuning. During training, we freeze the DINO image encoder, which shows excellent visual representation capacity, and only optimize the LoRA layers and depth decoder to integrate features from the surgical scene.

Results Our model is extensively validated on a MICCAI challenge dataset of SCARED, which is collected from da Vinci Xi endoscope surgery. We empirically show that Surgical-DINO significantly outperforms all the state-of-the-art models in endoscopic depth estimation tasks. The analysis with ablation studies has shown evidence of the remarkable effect of our LoRA layers and adaptation.

Conclusion Surgical-DINO shed some light on the successful adaptation of the foundation models into the surgical domain for depth estimation. There is clear evidence in the results that zero-shot prediction on pre-trained weights in computer vision datasets or naive fine-tuning is not sufficient to use the foundation model in the surgical domain directly.

Keywords Surgical scene understanding · Foundation models · Depth estimation · Adapter learning

Beilei Cui and Mobarakol Islam have contributed equally to this work.

✉ Hongliang Ren
hlren@ee.cuhk.edu.hk

Beilei Cui
beileicui@link.cuhk.edu.hk

Mobarakol Islam
mobarakol.islam@ucl.ac.uk

Long Bai
b.long@link.cuhk.edu.hk

- ¹ The Chinese University of Hong Kong, Hong Kong, China
- ² Wellcome/EPSCRC Centre for Interventional and Surgical Sciences (WEISS), University College London, London, UK
- ³ Department of BME, National University of Singapore, Singapore, Singapore

Introduction

3D scene reconstruction in endoscopic surgery has a significant impact on the development of automated surgery and promotes the advancement of various downstream applications such as surgical navigation, depth perception, augmented reality [1–3]. However, there are still many unresolved challenges in dense depth estimation tasks within endoscopic scenes. The variability of soft tissues and occlusion by surgical tools in the surgical environment poses high demands on the model's ability to reconstruct dynamic depth maps [4]. Recent methods have focused on utilizing binocular information to obtain disparity maps and reconstruct depth information [1, 4]. However, apart from the da Vinci surgical robot system, most endoscopic surgical robot systems only consist of a monocular camera, which is a more cost-effective and easily implementable hardware solution. Therefore, pre-

cise depth estimation tasks based on monocular endoscopy are still an area that requires further exploration.

Recently, foundation models have become one of the most popular terms in the field of deep learning [5, 6]. Thanks to their large number of model parameters, foundation models have the ability to build long-term memory of massive training data, achieving state-of-the-art performance on various downstream tasks involving vision, text and multimodal inputs. However, when encountering domain-specific scenarios such as surgical scenes, the predictive ability of foundation models tends to decline significantly [7]. Due to the limited availability of annotated data in medical scenes and insufficient computational resources, training a medical-specific foundation model from scratch poses various challenges. Therefore, there has been extensive discussion on adapting existing foundation models to different sub-domains, maximizing the utilization of existing model parameters, and fine-tuning foundation models for target application scenarios based on limited computational resources [7–9]. Chen et al. [8] constructed their adapter using two MLP layers and an activation function without inputting any prompt for fine-tuning the Segment Anything (SAM) model. On the other hand, Wu et al. [9] used a simple pixel classifier as a self-prompt to achieve zero-shot segmentation based on SAM. However, the adapter layer shall slow down inference speed, and prompts cannot be directly optimized through training. Therefore, we have designed our adaptation solution based on Low-Rank Adaptation (LoRA) [10]. LoRA adds a bypass next to the original foundation model, which performs a dimensionality reduction and then an elevation operation to simulate the intrinsic rank. When deployed in a production environment, LoRA can be introduced without introducing inference delays, and only the pre-trained model parameters need to be combined with the LoRA parameters. Therefore, LoRA can serve as an efficient adaption tool in real-world applications of foundation models.

Additionally, current works on fine-tuning vision foundation models to the medical domain have focused on common tasks such as segmentation and detection, with limited exploration in pixel-wise regression tasks like depth estimation. In this case, supervised training paradigms for visual foundation models are typically applied to common semantic understanding tasks and may not be suitable for our needs. Therefore, we have chosen DINOv2 [6] as the starting point for our study in this paper. DINOv2 is a self-supervised trained foundation model for multiple vision tasks. The self-supervised training paradigm enables DINOv2 to effectively learn unified visual features, thereby requiring only customized decoders to adapt DINOv2 to various downstream visual tasks including depth estimation. Therefore, we aim to explore the fine-tuning of the DINOv2 encoder to fully utilize the pre-trained extensive parameters and benefit downstream

depth estimation tasks in the surgical domain. Specifically, our key contributions and findings are:

- We firstly extend the foundation model in computer vision, DINOv2, to explore its capability on medical image depth estimation problems.
- We present an adaptation and fine-tuning strategy of DINOv2 based on the Low-Rank Adaptation technique with low additional training costs toward the surgical image domain.
- Our method, Surgical-DINO, is validated on two publicly available datasets and obtained superior performance over other state-of-the-art depth estimation methods for surgical images. We also investigate that the zero-shot foundation model is not yet ready for use in surgical applications, and LoRA adaptation is crucial, which outperformed naive fine-tuning.

Methodology

Preliminaries

DINOv2

Learning pre-trained representations without regard to specific tasks has been proven extremely effective in Natural Language Processing (NLP) [11]. One can use features from these pre-trained representations without fine-tuning for downstream tasks and obtain significantly better performances than those task-specific models. Oquab et al. [6] developed a similar "foundation" model, named DINOv2, in computer vision where vision features at both image level and pixel level generated from it can work without any task limitation. They proposed an automatic pipeline to build a large, curated and dedicated image dataset and an unsupervised learning method to learn robust vision features. A ViT model [12] with 1B parameters was trained in a discriminative self-supervised training manner and distilled into a series of smaller models that were evaluated to have surpassing ability against the best available all-purpose features on most of the benchmarks at image and pixel levels. Depth estimation task was also tested as a classical dense prediction task in computer vision by training a simple depth decoder head following DINOv2 and gained excellent performance in the general computer vision realm. The huge domain gap between medical and natural images may impede the utilization of such a foundation model; thus, we first attempt to develop a simple but effective adaptation method to exploit DINOv2 for the surgical domain.

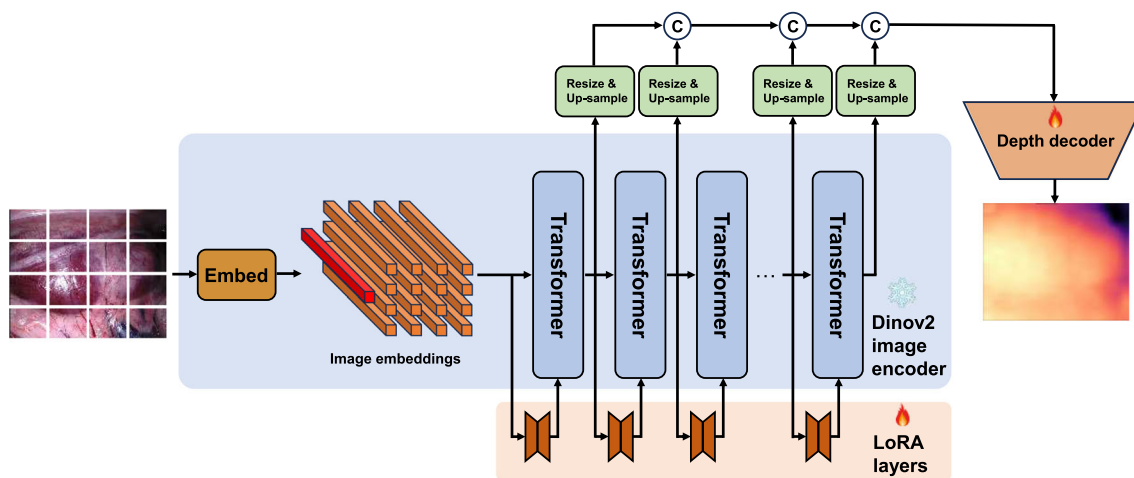


Fig. 1 The proposed Surgical-DINO framework. The input image is transformed into tokens by extracting scaled-down patches followed by a linear projection. A positional embedding and a patch-independent class token (red) are used to augment the embedding subsequently. We freeze the image encoder and add trainable LoRA layers to fine-tune

the model. We extract tokens from different layers, then up-sample and concatenate them to form the embedding features. Another trainable decode head is used on top of the frozen model to estimate the final depth

LoRA

Low-Rank Adaptation (LoRA) was first proposed in [10] to fine-tune large-scale foundation models in NLP to downstream tasks. It was inspired by the low “intrinsic dimension” of the pre-trained large model that random projection to a smaller subspace does not affect its ability to learn effectively. By injecting trainable rank decomposition matrices into each layer of the Transformer architecture and freezing the pre-trained model weights, LoRA significantly reduces the amount of trainable parameters for downstream tasks. To be specific, for a pre-trained weight matrix $W_0 \in \mathbb{R}^{d \times k}$, LoRA utilizes a low-rank decomposition to restrict its update by $W_0 + \Delta W = W_0 + BA$ where $B \in \mathbb{R}^{d \times r}$, $A \in \mathbb{R}^{r \times k}$ with the rank $r \ll \min(d, k)$. W_0 does not receive gradient updates during the training process, while only A and B contain trainable parameters. The modified forward pass is then described as:

$$h = W_0x + \Delta Wx = W_0x + BAx. \tag{1}$$

This implementation can significantly reduce the memory and storage usage for training thus very suitable for fine-tuning large-scale foundational models to downstream tasks.

Surgical-DINO

As illustrated in Fig. 1, the architecture of our proposed Surgical-DINO depth estimation framework inherits from DINOv2. Given a surgical image $x \in \mathbb{R}^{H \times W \times C}$ with spatial resolution $H \times W$ and channels C , we aim to predict its depth

map $\hat{D} \in H \times W$ as close to ground truth depth as possible. DINOv2 serves as an image encoder where images are first split into patches of size p^2 and then flattened with linear projection. A positional embedding is augmented for the tokens and another learnable class token is added which aggregates the global image information for subsequent missions. The image embeddings will then go through a series of Transformer blocks to generate new token representations. All parameters in the DINOv2 image encoder are frozen during training, and we added additional LoRA layers to each Transformer block to capture the learnable information. These side LoRA layers, as illustrated in the previous section, compress the Transformer vision features to the low rank space and then re-project back to match the output features’ channels in the frozen transformer blocks. LoRA layers in each Transformer block work independently and do not share weights. Several intermediate and the final output token representations will be resized and bi-linearly upsampled by a factor of 4 first and then concatenated to output the overall feature representation. A simple trainable depth decoder head is utilized at the end to predict the depth map.

LoRA layers

Different from fine-tuning the whole model, freezing the model and adding trainable LoRA layers will largely reduce the required memory and computation resources for training and also benefit conveniently deploying the model. The LoRA design in Surgical-DINO is presented in Fig. 2. We followed [13] where the low-rank approximation is only applied for q and v projection layers to avoid excessive influence on

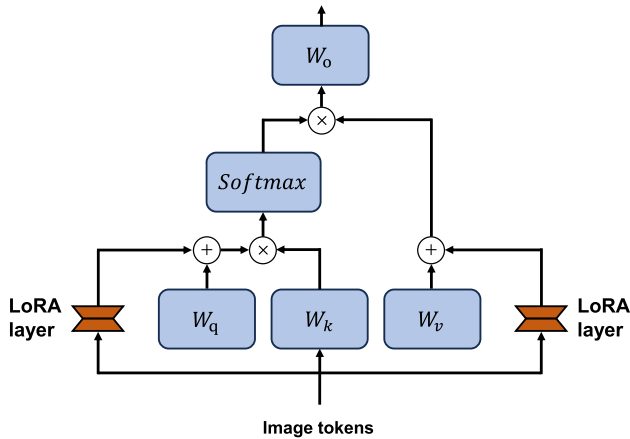


Fig. 2 The LoRA design in Surgical-DINO. We apply LoRA only to q and v projection layers in each transformer block. W_q , W_k , W_v and W_o denote the projection layer of q , k , v and o , respectively

attention scores. With the aforementioned fundamental formulation of LoRA, for an encoded token embedding x , the processing of q , k and v projection layers within a multi-head self-attention block will become:

$$\begin{aligned} Q &= \hat{W}_q a = W_q a + B_q A_q a, \\ K &= W_k a, \\ V &= \hat{W}_v a = W_v a + B_v A_v a, \end{aligned} \quad (2)$$

where W_q , W_k and W_v are frozen projection layers for q , k and v ; A_q , B_q , A_v and B_v are trainable LoRA layers. The self-attention mechanism remains unchanged that described by:

$$\text{Att}(Q, K, V) = \text{Softmax} \left(\frac{QK^T}{\sqrt{C_{\text{out}}}} + B \right) V \quad (3)$$

where C_{out} denotes the numbers of output tokens.

Network architecture

Image Encoder. The image is first separated into non-overlapping patches and then projected to image embeddings with the Embedding process. The image embeddings are a set of $t^0 = \{t_0^0, \dots, t_{N_p}^0\}$, $t_n^0 \in \mathbb{R}^D$ tokens, where p is the patch size, $N_p = \frac{HW}{p^2}$, t_0 is the class token, and D represents the feature dimensions of each token. L Transformers are then used to transform the image tokens into feature representations t^l where l denotes the output of l -th Transformer block. We utilized the pre-trained ViT-Base model from DINOv2 as our image encoder with 12 Transformer blocks and a feature dimension of 784.

Depth Decoder. We first extract the layers from $l = \{3, 6, 9, 12\}$, unflatten them to fit the patch resolution and

up-sample tokens by a factor of 4 to increase the resolution. We treat depth prediction as a classification problem by dividing the depth range into 256 uniformly distributed bins with a linear layer to predict the depth. The predicted map is scaled to align the input resolution eventually.

Loss functions

Surgical-DINO utilizes Scale-invariant depth loss [14] and Gradient loss [15] as the supervision constraints for the fine-tuning process. They can be described by:

$$\begin{aligned} \mathcal{L}_{\text{pixel}} &= \lambda_1 \sqrt{\frac{1}{n} \sum_i (g_i)^2 - \frac{\lambda_2}{n^2} \left(\sum_i g_i \right)^2} \\ \mathcal{L}_{\text{grad}} &= \lambda_3 \frac{1}{n} \sum_k \sum_i \left(|\nabla_x g_i^k| + |\nabla_y g_i^k| \right) \end{aligned} \quad (4)$$

where n denotes the number of valid pixels; $g_i^k = \log \tilde{d}_i^k - \log d_i^k$ is the value of the log-depth difference map at position i and scale k . $\mathcal{L}_{\text{pixel}}$ guides the network to predict truth depth, while $\mathcal{L}_{\text{grad}}$ encourages the network to predict smoother gradient changes. The final loss is then described as:

$$\mathcal{L} = \mathcal{L}_{\text{pixel}} + \mathcal{L}_{\text{grad}}. \quad (5)$$

Experiment

Dataset

SCARED¹ dataset is collected with a da Vinci Xi endoscope from fresh porcine cadaver abdominal anatomy and contains 35 endoscopic videos with 22,950 frames. A projector is used to obtain high-quality depth maps of the scene. Each video has ground truth depth and ego-motion, while we only used depth to evaluate our method. We followed the split scheme in [16] where the SCARED dataset is split into 15,351, 1705 and 551 frames for the training, validation and test sets, respectively.

Hamlyn² is a laparoscopic and endoscopic video dataset taken from various surgical procedures with challenging in vivo scenes. We followed the selection in [17] with 21 videos for validation.

Implementation details

The framework is implemented with PyTorch on NVIDIA RTX 3090 GPU. We adopt the AdamW [18] optimizer with

¹ <https://endovissub2019-scared.grand-challenge.org/>.

² <https://hamlyn.doc.ic.ac.uk/vision/>.

Table 1 Quantitative depth comparison on the SCARED dataset of SOTA depth estimation methods

Method	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	δ ↑
SfMLearner [19]	0.079	0.879	6.896	0.110	0.947
Fang et al. [20]	0.078	0.794	6.794	0.109	0.946
Defeat-Net [21]	0.077	0.792	6.688	0.108	0.941
SC-SfMLearner [22]	0.068	0.645	5.988	0.097	0.957
Monodepth2 [23]	0.071	0.590	5.606	0.094	0.953
Endo-SfM [24]	0.062	0.606	5.726	0.093	0.957
AF-SfMLearner [16]	<i>0.059</i>	0.435	4.925	0.082	<i>0.974</i>
DINOv2 [6] (zero-shot)	0.088	0.963	7.447	0.120	0.933
DINOv2 [6] (fine-tuned)	0.060	0.459	<i>4.692</i>	<i>0.081</i>	0.963
Surgical-DINO SSL (Ours)	<i>0.059</i>	<i>0.427</i>	4.904	<i>0.081</i>	0.974
Surgical-DINO (Ours)	0.053	0.377	4.296	0.074	0.975

The best results are in bold. The second-best results are in italic

Table 2 Quantitative depth comparison on Hamlyn dataset

Method	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	δ ↑
Endo-depth-and-motion [17]	0.185	5.424	16.100	0.225	0.732
AF-SfMLearner [16]	<i>0.168</i>	<i>4.440</i>	<i>13.870</i>	<i>0.204</i>	<i>0.770</i>
Surgical-DINO (ours)	0.146	3.216	11.974	0.178	0.801

The best results are in bold. The second-best results are in italic

an initial learning rate of 1×10^{-5} and weight decay of 1×10^{-4} . The batch size is set to 8 with 50 epochs in total. We can achieve our evaluation results with the following weights set: $\lambda_1 = 1.0$, $\lambda_2 = 0.85$, $\lambda_3 = 0.5$. The images are resized to 224×224 pixels. We also trained our proposed model in a Self-Supervised Learning (SSL) manner with the baseline of AF-SfMLearner [16]. We replace the encoder in AF-SfMLearner with Surgical-DINO and resize the image to 224×224 pixels to fit the patch size of DINOv2.

Performance metrics

We evaluate our method with five common metrics used in depth estimation problems: Abs Rel, Sq Rel, RMSE, RMSE log and δ in which lower is better for the first four metrics and larger is better for the last one. During evaluation, we re-scale the predicted depth map by a median scaling method introduced by SfM-Learner [19], which can be expressed by

$$\mathbf{D}_{\text{scaled}} = (\mathbf{D}_{\text{pred}} * (\text{median}(\mathbf{D}_{\text{gt}}) / \text{median}(\mathbf{D}_{\text{pred}}))) \quad (6)$$

We capped the depth map at 150mm which can cover almost all depth values.

Results

Quantitative results on SCARED. We compared our proposed method with several SOTA self-supervised methods [16, 19–24] as well as zero-shot, self-supervised and supervised method, and the results are shown in Table 1. All

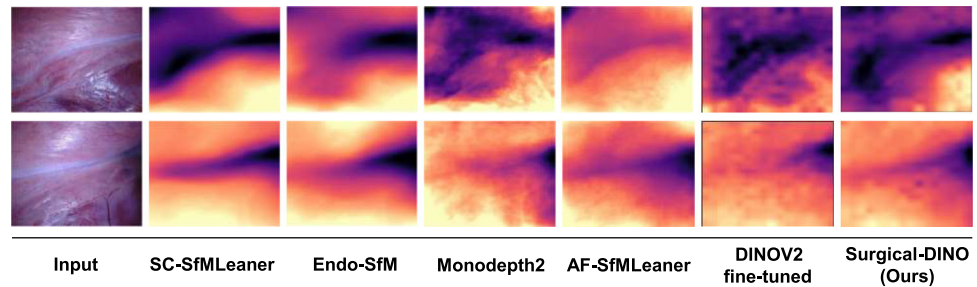
of these baseline methods were reproduced with the original implementation under the same dataset splits mentioned above. The zero-shot performance of pre-trained DINOv2 is evaluated on model size ViT-Base with a same depth decoder head fine-tuned on NYU Depth V2 [25]. Our method obtained superior performances in all the evaluation metrics compared to all of the methods. It is worth noting that the zero-shot performance of DINOv2 has the worst results indicating that vision features and depth decoder that are highly effective in natural images are unsuitable for medical images due to the large domain gap. While the fine-tuned DINOv2 exceeds other SOTA self-supervised methods in RMSE and RMSE log, it did not gain better performance in the other three metrics proving its prediction to have more large depth errors. Only fine-tuning a depth decoder head is not enough to transfer the vision features to geometric relations within medical images. With the adaptation method of LoRA, the network is able to learn medical domain-specific vision features and relate them with depth information, thus resulting in an improvement in the estimation accuracy.

Quantitative results on Hamlyn. We made zero-shot validation for our model trained on SCARED in Hamlyn dataset without any fine-tuning. For comparison, we zero-shot validate AF-SfMLearner with their best model and obtain the results of Endo-Depth-and-Motion [17] by averaging the 21-fold cross-validation results trained on Hamlyn. As presented in Table 2, our method achieves superior performance against other methods, unveiling the good generalization ability across different cameras and surgical scenes.

Table 3 Comparison of encoder parameters, trainable parameters, trainable parameters' ratio and full model inference speed

Method	Params. (M) ↓	trainable Params. (M) ↓	trainable ratio (%) ↓	Speed (ms) ↓
AF-SfMLearner [16]	11.68	11.68	100.00	9.9
Surgical-DINO (Ours)	86.72	0.14	0.17	18.2

The best results are in bold

Fig. 3 Qualitative depth comparison on the SCARED dataset**Table 4** Ablation study on the rank size on the LoRA layer

Rank size	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	δ ↑
1	0.058	0.389	4.513	0.081	0.962
4	0.053	0.375	4.296	0.074	0.975
8	0.053	0.376	4.324	0.074	0.974
16	0.053	0.377	4.325	0.073	0.974

The best results are in bold

Model complexity and speed evaluation. The proposed model's parameters, trainable parameters, trainable parameters ratio and inference speed are evaluated on an NVIDIA RTX 3090 GPU compared to AF-SfMLearner. Table 3 shows that while Surgical-DINO has a larger amount of parameters, only a very small part of parameters are trainable making it faster to train and converge. The inference speed of Surgical-DINO is slower than AF-SfMLearner, but still in an acceptable range for real-time applications.

Qualitative results. We also show some qualitative results in Fig. 3. Our method can depict anatomical structure well compared to other methods. Nevertheless, the qualitative results of our proposed Surgical-DINO also have drawbacks like lack of continuity which can motivate future modification direction.

Ablation studies

Effects of the rank size on the LoRA layer. A set of comparative experiments is performed to evaluate the effects of

rank size on the LoRA layer. We evaluated four different sizes of rank of the LoRA layer, and the results are shown in Table 4. We notice that the performance of Surgical-DINO will increase with the increase of rank size within a certain low range and start to drop when the rank size exceeds a certain value. This phenomenon implies that despite being designed to utilize low-rank decomposition to make the approximation, LoRA still requires certain training parameters to fit downstream tasks. However, too many trainable parameters may mislead the original weights resulting in performance degradation.

Effects of the size of pre-trained foundation model. DINOv2 published four pre-trained ViT foundation models and named them by their size. Table 5 presents the ablation study to investigate the effect of the size of the pre-trained foundation model. We discover that the performance increases with the increase of the pre-trained model size. Larger models inherently have better integration and generalization ability of vision features, thus better fitting downstream tasks. But larger models are also accompanied by larger memory

Table 5 Ablation study on the size of pre-trained foundation model

Model size	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	δ ↑
ViT-small	0.055	0.416	4.513	0.075	0.971
ViT-base	0.053	0.377	4.296	0.074	0.975
ViT-large	0.051	0.363	4.256	0.070	0.979
ViT-giant	0.050	0.342	4.120	0.069	0.980

occupancy and training costs, so we chose ViT-Base for our depth estimation method in consideration of the compromise between performance and cost.

Conclusions

Depth estimation is a vital task in robotic surgery and benefits many downstream tasks like surgical navigation and 3D reconstruction. Vision Foundation model that captures universal vision features has been proven to be both effective and convenient in many vision tasks but yet needs more exploration in the surgical domain. We have presented Surgical-DINO, an adapter learning method that utilizes DINOv2, a vision foundation model, for surgical scene depth estimation. We design LoRA layers to fine-tune the network with a small number of additional parameters to adapt to the surgical domain. Experiments have been made on a publicly available dataset and demonstrate the superior performance of the proposed Surgical-DINO. We first explore the direction of deploying the vision foundation model to surgical depth estimation tasks and reveal its enormous potential. Future works could explore the foundation model in a supervised, self-supervised and unsupervised manner to investigate the robustness and reliability in the surgical domain.

Funding This work was supported by Hong Kong Research Grants Council (RGC) Collaborative Research Fund (C4026-21G), General Research Fund (GRF 14211420 & 14203323), Shenzhen-Hong Kong-Macau Technology Research Programme (Type C) STIC Grant SGDX20210823103535014 (202108233000303).

Code availability The source code is available at <https://github.com/BeileiCui/SurgicalDINO>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

Informed consent This article does not contain patient data.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Zha R, Cheng X, Li H, Harandi M, Ge Z (2023) Endosurf: neural surface reconstruction of deformable tissues with stereo endoscope videos. International conference on medical image computing and computer-assisted intervention. Springer, Berlin, pp 13–23
- Liu X, Sinha A, Ishii M, Hager GD, Reiter A, Taylor RH, Unberath M (2019) Dense depth estimation in monocular endoscopy with self-supervised learning methods. *IEEE Trans Med Imaging* 39(5):1438–1447
- Wei X, Wang Y, Ge L, Peng B, He Q, Wang R, Huang L, Xu Y, Luo J (2022) Unsupervised convolutional neural network for motion estimation in ultrasound elastography. *IEEE Trans Ultrason Ferroelectr Freq Control* 69(7):2236–2247
- Wang Y, Long Y, Fan SH, Dou Q (2022) Neural rendering for stereo 3d reconstruction of deformable tissues in robotic surgery. International conference on medical image computing and computer-assisted intervention. Springer, Berlin, pp 431–441
- Kirillov A, Mintun E, Ravi N, Mao H, Rolland C, Gustafson L, Xiao T, Whitehead S, Berg AC, Lo W-Y, Dollár P, Girshick R (2023) Segment anything. *arXiv preprint arXiv:2304.02643*
- Oquab M, Darcet T, Moutakanni T, Vo H, Szafraniec M, Khalidov V, Fernandez P, Haziza D, Massa F, El-Nouby A, Assran M, Ballas N, Galuba W, Howes R, Huang P-Y, Li S-W, Misra I, Rabbat M, Sharma V, Synnaeve G, Xu H, Jegou H, Mairal J, Labatut P, Joulin A, Bojanowski P (2023) DINOv2: learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*
- Wang A, Islam M, Xu M, Zhang Y, Ren H (2023) Sam meets robotic surgery: an empirical study on generalization, robustness and adaptation. *arXiv preprint arXiv:2308.07156*
- Chen T, Zhu L, Ding C, Cao R, Zhang S, Wang Y, Li Z, Sun L, Mao P, Zang Y (2023) Sam fails to segment anything?—sam-adapter: adapting sam in underperformed scenes: Camouflage, shadow, and more. *arXiv preprint arXiv:2304.09148*
- Wu Q, Zhang Y, Elbatel M (2023) Self-prompting large vision models for few-shot medical image segmentation. MICCAI workshop on domain adaptation and representation transfer. Springer, Berlin, pp 156–167
- Hu EJ, yelong shen Wallis P, Allen-Zhu Z, Li Y, Wang S, Wang L, Chen W (2022) LoRA: low-rank adaptation of large language models. In: International conference on learning representations
- Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M, Zhou Y, Li W, Liu PJ (2020) Exploring the limits of transfer learning with a unified text-to-text transformer. *J Mach Learn Res* 21(1):5485–5551
- Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, Uszkoreit J, Houlsby N (2021) An image is worth 16x16 words: transformers for image recognition at scale. In: International conference on learning representations
- Zhang K, Liu D (2023) Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785*
- Bhat SF, Alhashim I, Wonka P (2021) Adabins: depth estimation using adaptive bins. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 4009–4018
- Li Z, Snavely N (2018) Megadepth: learning single-view depth prediction from internet photos. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 2041–2050
- Shao S, Pei Z, Chen W, Zhu W, Wu X, Sun D, Zhang B (2022) Self-supervised monocular depth and ego-motion estimation in endoscopy: appearance flow to the rescue. *Med Image Anal* 77:102338
- Recasens D, Lamarca J, Fácil JM, Montiel J, Civera J (2021) Endo-depth-and-motion: reconstruction and tracking in endo-

- scopic videos using depth networks and photometric constraints. *IEEE Robot Autom Lett* 6(4):7225–7232
18. Loshchilov I, Hutter F (2019) Decoupled weight decay regularization. In: *International conference on learning representations*
 19. Zhou T, Brown M, Snavely N, Lowe DG (2017) Unsupervised learning of depth and ego-motion from video. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 1851–1858
 20. Fang Z, Chen X, Chen Y, Gool LV (2020) Towards good practice for CNN-based monocular depth estimation. In: *Proceedings of the IEEE winter conference on applications of computer vision*, pp 1091–1100
 21. Spencer J, Bowden R, Hadfield S (2020) Defeat-net: general monocular depth via simultaneous unsupervised representation learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 14402–14413
 22. Bian J, Li Z, Wang N, Zhan H, Shen C, Cheng M-M, Reid I (2019) Unsupervised scale-consistent depth and ego-motion learning from monocular video. *Adv Neural Inf Process Syst* 32
 23. Godard C, Mac Aodha O, Firman M, Brostow GJ (2019) Digging into self-supervised monocular depth estimation. In: *Proceedings of the IEEE international conference on computer vision*, pp 3828–3838
 24. Ozyoruk KB, Gokceler GI, Bobrow TL, Coskun G, Incetan K, Almalioglu Y, Mahmood F, Curto E, Perdigoto L, Oliveira M, Sahin H, Araujo H, Alexandrino H, Durr NJ, Gibert HB, Mehmet T (2021) Endoslam dataset and an unsupervised monocular visual odometry and depth estimation approach for endoscopic videos. *Med Image Anal* 71:102058
 25. Silberman N, Hoiem D, Kohli P, Fergus R (2012) Indoor segmentation and support inference from RGBD images. In: *Computer vision—ECCV 2012: 12th European conference on computer vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*, pp 746–760. Springer

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.