

EndoDDC: Learning Sparse to Dense Reconstruction for Endoscopic Robotic Navigation via Diffusion Depth Completion

Yinheng Lin^{1†}, Yiming Huang^{1†}, Beilei Cui¹, Long Bai¹, Huxin Gao¹, Hongliang Ren^{1,2} and Jiewen Lai^{1*}

Abstract—Accurate depth estimation plays a critical role in the navigation of endoscopic surgical robots, forming the foundation for 3D reconstruction and safe instrument guidance. Fine-tuning pretrained models heavily relies on endoscopic surgical datasets with precise depth annotations. While existing self-supervised depth estimation techniques eliminate the need for accurate depth annotations, their performance degrades in environments with weak textures and variable lighting, leading to sparse reconstruction with invalid depth estimation. Depth completion using sparse depth maps can mitigate these issues and improve accuracy. Despite the advances in depth completion techniques in general fields, their application in endoscopy remains limited. To overcome these limitations, we propose EndoDDC, an endoscopy depth completion method that integrates images, sparse depth information with depth gradient features, and optimizes depth maps through a diffusion model, addressing the issues of weak texture and light reflection in endoscopic environments. Extensive experiments on two publicly available endoscopy datasets show that our approach outperforms state-of-the-art models in both depth accuracy and robustness. This demonstrates the potential of our method to reduce visual errors in complex endoscopic environments. Our code will be released at <https://github.com/yinheng-lin/EndoDDC>.

I. INTRODUCTION

Accurate navigation of endoscopic surgical robots is fundamental to minimally invasive surgery, enabling surgeons to achieve improved precision, flexibility, and patient outcomes [1]. A critical challenge for navigation of these systems is accurate depth estimation. These robots acquire depth maps from 2D images and convert them into 3D representations, which is essential for improving spatial awareness, enabling autonomous path planning, and supporting real-time navigational decisions within complex anatomical structures [2]. Despite their promising potential, accurate depth perception remains a primary challenge for the navigation of endoscopic surgical robot systems. One common approach involves fine-tuning state-of-the-art (SOTA) models

This work was supported in part by the Hong Kong Innovation and Technology Fund (ITF) (Grant No. MHP/185/24); in part by the National Key Research and Development Program of China (Grant No. 2024YFE0216200); in part by the Guangdong Basic and Applied Basic Research Foundation (Grant No. 2025A1515011594); in part by the National Natural Science Foundation of China (Grant No. 62403402) and NSFC Distinguished Young Scientists Fund – Category A (Grant No. T252500134); and in part by the Hong Kong Research Grants Council (Grant Nos. C4026-21GF, R4020-22, 14200425, 14206125, 14204524, 14203323) (*Corresponding author: Jiewen Lai)

¹All authors are with the Department of Electronic Engineering, The Chinese University of Hong Kong, Hong Kong, China. {linyinheng, yhuangdl, beileicui, b.long}@link.cuhk.edu.hk, {hxgao, hlren, jwlai}@ee.cuhk.edu.hk

²H. Ren is also with the Shenzhen Loop Area Institute, Shenzhen, China.

[†]These authors contributed equally to this work.

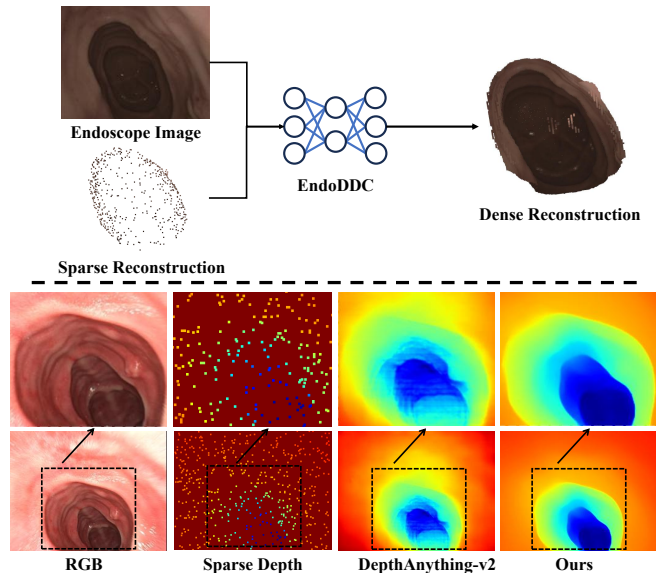


Fig. 1. Compared with fine-tuning the SOTA foundational model, our method generates more robust and accurate depth from sparse depth and RGB input, achieving superior sparse to dense reconstruction.

on surgical datasets [3]. However, this requires dense depth ground truth that is difficult to acquire. The endoscopic environment, characterized by textureless tissue surfaces and specular reflections, fundamentally limits data acquisition. Consequently, depth sensors such as Time-of-Flight (ToF) [4] and stereo endoscopes [5] provide only sparse and incomplete depth measurements. While these sparse points are accurate, they are insufficient for the model fine-tuning to reconstruct a complete and dense 3D scene.

Self-supervised regression methods, independent of datasets, provide a promising alternative [6]. However, monocular depth estimation using single RGB images suffers from scale ambiguity and struggles to produce geometrically accurate results, particularly in weakly textured endoscopic environments. Self-supervised methods using video sequences depend on significant camera motion to establish geometric constraints, making them unsuitable for endoscopic scenarios, where minimal camera movement complicates depth and pose estimation [7].

While self-supervised depth estimation methods suffer from scale ambiguity issues and fail to generate precise depth maps, depth completion with accurate sensor depth becomes a more viable alternative. By leveraging accurate sparse depth points as a geometric prior, depth completion methods address challenges inherent in fine-tuning methods

and self-supervised learning, generating dense depth maps by combining sparse depth data with RGB images [8]. While depth completion has seen success in fields such as autonomous driving [9], due to the weak texture and multiple light reflections in endoscopic environments, its application in endoscopic robotic navigation remains unexplored. Additionally, the sparse depth values obtained via depth sensors are difficult to utilize for model fine-tuning.

To tackle these challenges, we introduce EndoDDC, a novel depth completion method tailored to the navigation of endoscopic surgical robots. Our approach integrates depth gradient information and introduces a diffusion model to optimize depth maps, addressing the issues of weak texture and light reflection in endoscopic environments. The method consists of a feature extraction network, a depth fusion module for iterative processing of depth values and gradients, a conditional depth diffusion model for refining the predicted depth map, and an upsampling module for producing high-resolution results. As shown in Fig. 1, our method outperforms fine-tuning the current SOTA model, DepthAnything-V2 [10], leveraging sparse depth information to compensate for missing gastrointestinal features, significantly improving the geometric and textural details of the intestinal wall.

Our key contributions include the following:

- We propose EndoDDC, a depth completion pipeline that leverages sparse depth and RGB input for sparse to dense reconstruction, addressing challenges posed by fine-tuning and self-supervised learning methods.
- We propose a multi-scale feature extraction and depth gradient fusion module to provide geometry guidance for sparse to dense endoscopic reconstruction.
- We propose a depth diffusion strategy with the depth gradient conditioned diffusion model for the endoscopic depth completion task, iteratively optimizing the depth values with coarse depth and depth gradient features.
- We conduct exhaustive experiments on two public surgical datasets, showing that our model outperforms current SOTA methods by providing more robust sparse to dense reconstruction for endoscopic robotic navigation.

II. RELATED WORK

A. Depth Estimation in Endoscopic Surgery

Recent advances in endoscopic surgery have leveraged the powerful prior knowledge of foundational models [10]–[13] by fine-tuning them on surgical data, enabling precise three-dimensional perception in complex surgical environments. However, the success of this approach heavily depends on the availability of domain-specific datasets with accurate ground truth depth. Due to restrictions related to safety, privacy, and professional regulations, obtaining precise large-scale endoscopic surgery videos with corresponding depth of ground truth remains a significant challenge [14], [15]. As a result, deep learning-based methods for endoscopic depth prediction are shifting toward self-supervised regression.

Simultaneous localization and mapping with multi-view learning, a fundamental paradigm in this domain, is based on

the Structure from Motion (SfM). Existing works [16]–[19] utilized representation such as ORB features [20] and 3D Gaussians [21], paving the way in learning-based endoscopic reconstruction. However, the limited range of motion in endoscopic cameras, confined within a cavity, often results in minimal pose changes. This makes the simultaneous estimation of depth and pose an ill-posed problem, diminishing the model’s accuracy and robustness [22]. To eliminate reliance on camera motion, self-supervised single-view methods, such as LightDepth [7], use unpaired single images for training, enabling robust depth estimation even in scenes with minimal or no camera movement. Nevertheless, since depth is inferred solely from single-frame data, the scale of depth values remains unknown. Moreover, variations in scene brightness can cause changes in the scale of each frame, making raw outputs unsuitable for clinical applications that require precise measurements. Therefore, additional mechanisms are necessary to restore the absolute scale [23].

B. Depth Completion

In conventional environments, early depth completion methods directly predicted depth values from images and sparse depth maps [8], [24]. These feedforward approaches struggle with inferring long-range depth due to their limited receptive fields. To improve output quality, post-processing optimization methods based on the spatial propagation network (SPN) mechanism [25] have been continuously proposed and refined. Specifically, CSPN [26] and its successor, CSPN++ [27], introduced convolutions with fixed and adaptive kernels, respectively, improving efficiency. NLSPN [28] utilizes deformable convolutions to predict non-local neighborhoods, while DySPN [29] predicts distinct propagation weights at each iteration, achieving superior performance with fewer iterations and smaller neighborhoods. Earlier architectures primarily processed two-dimensional features, but some methods, such as GuideFormer [30] and CompletionFormer [31], incorporate camera intrinsic parameters to handle three-dimensional information. Building on these advances, OGNI-DC [9] achieves more robust generalization by iteratively optimizing the depth gradient field.

However, most current methods focus on indoor and outdoor human activity scenarios, with limited efficient depth completion solutions for endoscopic surgery. Liu et al [32], proposed a depth completion method for intestinal endoscopy based on multi-scale confidence and self-attention mechanisms, but this method struggles with depth completion in regions with weak textures. To overcome these limitations, we propose a depth completion method based on depth gradients and diffusion optimization. With the physical nature of the diffusion model, our approach ensures precise depth reconstruction across various endoscopic datasets, advancing the dense 3D reconstruction techniques in surgical robotics.

III. METHODOLOGY

EndoDDC takes an RGB image $\mathbf{I} \in \mathbb{R}^{3 \times H \times W}$ and its corresponding sparse depth map $\mathbf{S} \in \mathbb{R}_+^{H \times W}$ as input, and

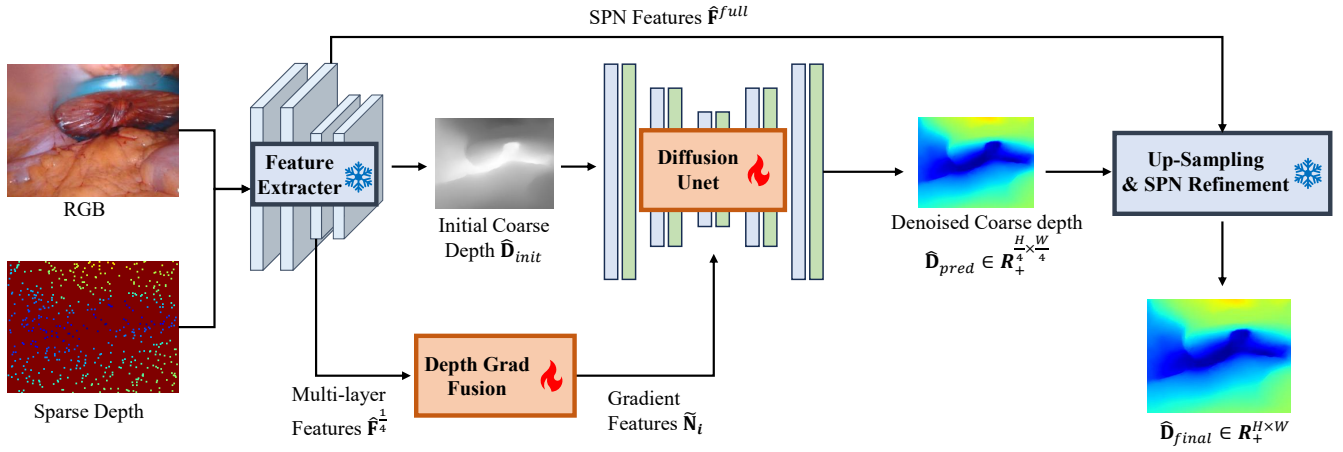


Fig. 2. **Overview of EndoDDC:** After feature extraction from the RGB image and sparse depth map, the Depth Grad Fusion module iteratively updates the state hidden network based on depth and gradient features. This output is then fed into the Depth Diffusion guidance model to conditionally optimize the initial depth. Finally, the optimized coarse depth map undergoes up-sampling and SPN refinement module to produce the final depth map.

ultimately outputs a dense depth map. The overall pipeline of our proposed framework is illustrated in Fig. 2.

A. Preliminary

Our depth diffusion module builds upon the Denoising Diffusion Implicit Model (DDIM) [33], a generative method that produces high-fidelity results through a deterministic denoising process. The DDIM framework consists of a forward diffusion process and a reverse denoising process.

Forward Diffusion. The forward process gradually adds Gaussian noise to an initial input \mathbf{x}_0 over T discrete timesteps. Following a Markov chain at timestep t : $q(x_{1:T} | x_0) = \prod_{t=1}^T q(x_t | x_{t-1})$, we have:

$$q(x_t | x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}\right) \quad (1)$$

where x_t indicates the noise data at time-step $t \in [0, T]$, β_t is the variance and $q(x_{1:T} | x_0)$ denotes the probability distribution of \mathbf{x}_T obtained after T -step. Based on the above equation, the relationship between \mathbf{x}_T and \mathbf{x}_0 is:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}_t \quad (2)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{i=0}^t \alpha_i$ and $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

Reverse Denoising. The reverse process recovers the original data \mathbf{x}_0 from a noisy input $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. At each step t , a network ϵ_θ first predicts the noise component $\boldsymbol{\epsilon}_t$ from the noisy input \mathbf{x}_T , then estimate the initial data as:

$$\hat{x}_0(x_t, t) = \frac{x_t - \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}_\theta}{\sqrt{\bar{\alpha}_t}}. \quad (3)$$

By the definition of DDIM, we have the denoised data:

$$x_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\hat{x}_0(x_t, t) + \sqrt{1 - \bar{\alpha}_{t-1}}\boldsymbol{\epsilon}_\theta(x_t, t) \quad (4)$$

This deterministic process is equivalent to solving a probability flow Ordinary Differential Equation (ODE) [34], [35]. Its iterative correction mechanism effectively rectifies errors caused by endoscopic challenges, like textureless surfaces or specular reflections, ensuring a globally coherent and physically plausible final depth map.

B. Feature Extraction and Depth Grad Fusion

In feature extraction, we first encode the input image \mathbf{I} and the sparse depth map \mathbf{S} through a series of convolutional layers. After concatenation, we use the pre-trained PVT [36] to extract both global and local information, generating feature maps at multiple scales. After encoding, the network enters the decoding stage, where it gradually restores spatial resolution through upsampling and skip connections. The network ultimately outputs the initial coarse depth prediction and features at two resolution scales: 1/4 resolution for intermediate depth predictions and full resolution for final depth enhancement with SPN:

$$\hat{\mathbf{D}}_{\text{init}}, \hat{\mathbf{F}}^{\text{full}}, \hat{\mathbf{F}}^{\frac{1}{4}} = \text{Backbone}(\mathbf{I}, \mathbf{S}) \quad (5)$$

Following [9], we use CompletionFormer [37] as the backbone of our feature extraction and initialize $\hat{\mathbf{G}}_0$ with 0.

Fig. 3(a) shows that the Depth Grad Fusion module iteratively processes depth maps and gradient maps using convolutional gated recurrent units (ConvGRU) [38] to obtain more precise feature networks containing depth and gradient information:

$$\Delta \hat{\mathbf{D}}_i, \Delta \hat{\mathbf{G}}_i, \tilde{\mathbf{N}}_i = \text{ConvGRU}\left(\tilde{\mathbf{N}}_{i-1}, \tilde{\mathbf{C}}_{i-1}, \hat{\mathbf{D}}_{i-1}, \hat{\mathbf{G}}_{i-1}\right) \quad (6)$$

For each step, this model takes hidden state net $\tilde{\mathbf{N}}_{i-1}$, context input $\tilde{\mathbf{C}}_{i-1}$, predicted depth gradient $\hat{\mathbf{G}}_{i-1}$, and predicted depth $\hat{\mathbf{D}}_{i-1}$ as input to generate the new hidden, variation of the context and depth gradient. The hidden state net $\tilde{\mathbf{N}}_0$ and init context input $\tilde{\mathbf{C}}_0$ are obtained from backbone features $\hat{\mathbf{F}}^{\frac{1}{4}}$. With the variation, we have the current depth $\hat{\mathbf{D}}_i = \Delta \hat{\mathbf{D}}_i + \hat{\mathbf{D}}_{i-1}$ and current depth gradient $\hat{\mathbf{G}}_i = \Delta \hat{\mathbf{G}}_i + \hat{\mathbf{G}}_{i-1}$. The final output of the hidden state net $\tilde{\mathbf{N}}_i$ is utilized as the depth diffusion condition, and the depth gradient $\hat{\mathbf{G}}_i$ serves as an intermediate for the depth gradient fusion optimization.

C. Depth Diffusion

The iterative denoising process of DDIM enforces a smooth path from noise to data, ensuring the global coher-

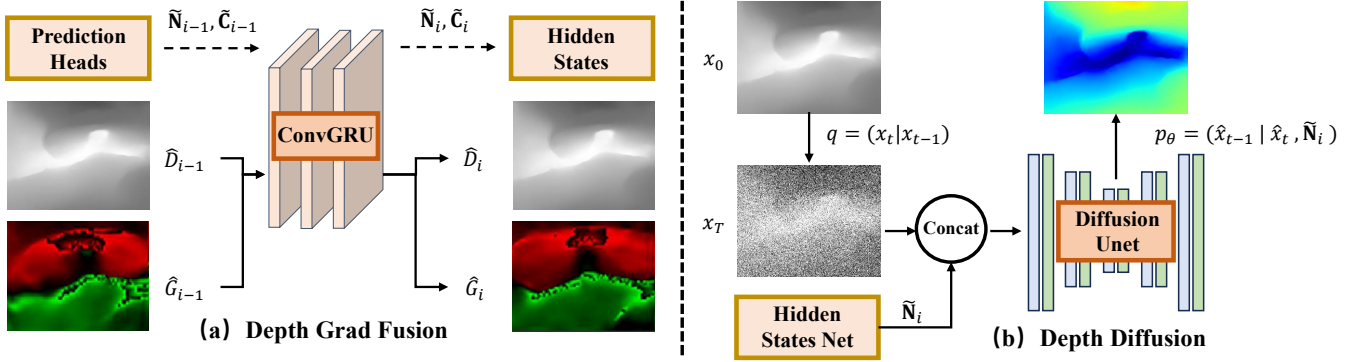


Fig. 3. **Details of the proposed model.** (a) Details of the Depth Grad Fusion module; (b) The architectural design of the Depth Diffusion module.

ence of the final depth map. However, while this mechanism is effective in maintaining overall plausibility, it lacks contextual initialization and explicit guidance to resolve the local ambiguities inherent in endoscopic imagery. To constrain the solution space and guide the generation process toward the correct geometry, we first utilize the initial coarse depth \hat{D}_{init} as the initialization x_t for the reverse denoising.

Then we introduce the explicit geometric diffusion condition by utilizing depth gradient features, which encode crucial information about local surface structure and orientation. As shown in Fig. 3(b), we denote the hidden states net containing depth and gradient information, extracted by a preceding network, as $\tilde{\mathbf{N}}_i \in \mathbb{R}^{C \times \frac{H}{4} \times \frac{W}{4}}$. And we distill this high-dimensional feature map into a single-channel guidance map $\tilde{\mathbf{N}}'_i \in \mathbb{R}^{1 \times \frac{H}{4} \times \frac{W}{4}}$ using a 1×1 convolution. This projection preserves the essential geometric cues while matching the dimensions of the depth map.

Then we incorporate this guidance by concatenating $\tilde{\mathbf{N}}'_i$ with the noisy depth map x_t along the channel axis. The combined tensor serves as the input to our noise prediction network. This modification explicitly informs the network about the underlying geometry at each denoising step, enabling it to make more accurate noise predictions. The conditional noise prediction is thus formulated as follows:

$$\epsilon_\theta = \epsilon_\theta(\text{concat}(x_t, \tilde{\mathbf{N}}'_i), t) \quad (7)$$

By training the network with this objective, our depth diffusion model learns to leverage geometric priors to effectively optimize the depth map. And the process of obtaining depth maps through iterative denoising, expressed in equation 4, can be rewritten as:

$$x_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \hat{x}_0(x_t, t) + \sqrt{1 - \bar{\alpha}_{t-1}} \epsilon_\theta(\text{concat}(x_t, \tilde{\mathbf{N}}'_i), t) \quad (8)$$

In the training phase, since $\bar{\alpha}_t$, t , $\tilde{\mathbf{N}}$ and x_t are known, we only need to train the network ϵ_θ to accurately predict the noise ϵ_t to get the corrected depth $\hat{D}_{pred} \in \mathbb{R}_+^{\frac{H}{4} \times \frac{W}{4}}$.

D. Depth Enhancement and Loss Function

Depth Up-Sampling and SPN Refinement. We use the convex combination method [38] to up-sample the $1/4$

resolution depth prediction map $\hat{D}_{pred} \in \mathbb{R}_+^{\frac{H}{4} \times \frac{W}{4}}$ to obtain the full resolution depth map $\hat{D}_{pred} \in \mathbb{R}_+^{H \times W}$, and then use a pre-trained spatial propagation network (SPN) [39] to augment the up-sampled depth map to obtain the final depth prediction map \hat{D}_{final} :

$$\hat{D}_{final} = \text{SPN}(\hat{D}_{pred}, \hat{\mathbf{F}}^{\text{full}}) \quad (9)$$

Depth Loss. During the training process, we combine L_1 and L_2 loss to supervise the final depth maps \hat{D}_{pred} and the up-sampled depth maps \hat{D}_{final} . Given the ground-truth depth map D , the loss function of depth \mathcal{L}_D is:

$$\mathcal{L}_D = \left\| \hat{D}_{pred} - D \right\|_2^2 + \left\| \hat{D}_{pred} - D \right\|_1 + \left\| \hat{D}_{final} - D \right\|_2^2 + \left\| \hat{D}_{final} - D \right\|_1 \quad (10)$$

Gradient Loss. In down-sampling, we compute the corresponding gradient G based on the provided ground-truth depth map D and supervise the depth gradients with L_1 :

$$\mathcal{L}_G = \gamma^{I-i} \left\| \hat{\mathbf{G}}_i - \mathbf{G} \right\|_1 \quad (11)$$

where I is the total number of iterations for convgru, i represents the current iteration, and loss decay γ is 0.9.

Depth Diffusion Loss. In the training of the depth diffusion model, we optimize $\epsilon_\theta(\text{concat}(x_t, \tilde{\mathbf{N}}'_i), t)$ to promote it close to the Gaussian noise, and this loss is supervised with L_2 and can be written as:

$$\mathcal{L}_{diff} = \left\| \epsilon_t - \epsilon_\theta(\text{concat}(x_t, \tilde{\mathbf{N}}'_i), t) \right\|_2 \quad (12)$$

IV. EXPERIMENTS

A. Datasets and Evaluation Metrics

C3VD (Colonoscopy 3D Video Dataset) [40] is a 3D colonoscopy video dataset acquired using high-definition clinical colonoscopes and high-fidelity colon models. The dataset comprises 10,015 annotated images and 22 registered video sequences, along with corresponding ground truth for depth, surface normals, optical flow, occlusion, six-degree-of-freedom (6-DoF) poses, coverage maps, and 3D models. In our experiments, we trained and tested C3VD using the dataset splits proposed in LightDepth [43].

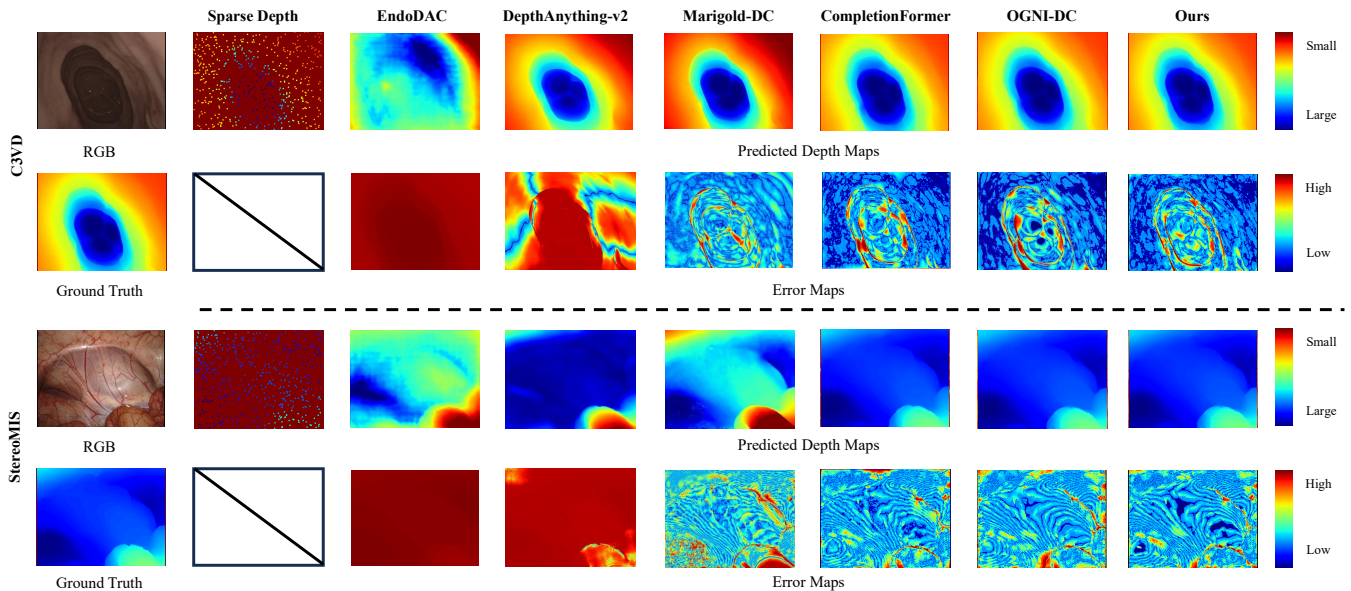


Fig. 4. **Qualitative comparison on C3VD and StereoMIS datasets.** We compare EndoDDC with SOTA depth estimate and depth completion methods; our method generates less error in terms of tissue details.

TABLE I
QUANTITATIVE DEPTH COMPARISON WITH SOTA METHODS ON C3VD [40] AND STEREOIMIS [41] DATASET.

Models	C3VD [40]				StereoMIS [41]			
	RMSE (mm) ↓	MAE (mm) ↓	REL (mm) ↓	δ ↑	RMSE (mm) ↓	MAE (mm) ↓	REL (mm) ↓	δ ↑
EndoDAC [22]	9.7476	7.5541	0.1081	0.9162	11.8435	8.7066	0.2084	0.6765
DepthAnything-v2 [10]	5.2202	3.6901	0.0671	0.9892	2.2465	1.63765	0.0277	0.8857
Marigold-DC [42]	0.8294	0.3275	0.0106	0.9985	2.6755	1.0395	0.0122	0.9968
CompletionFormer [31]	0.6875	0.2320	0.0071	0.9988	1.6536	0.6261	0.0069	0.9986
OGNI-DC [9]	0.6770	0.2283	0.0067	0.9988	1.5857	0.6114	0.0063	0.9986
Ours	0.6412	0.2104	0.0060	0.9990	1.4691	0.5515	0.0061	0.9988

StereoMIS [41] is a dataset designed for SLAM applications in endoscopic surgery. It was captured using the da Vinci Xi surgical robot to record stereo video streams and forward camera motions across three live porcine subjects. It encompasses diverse anatomical structures and challenging scenarios involving significant tissue deformation. We extracted images from the video frames at a frame ratio of 10:1 for its 11 scene videos. We used the images extracted from video frames P1 through P2_7 as the training set and the images extracted from P2_8 as the test set. In addition, since the dataset does not provide the ground truth of the paired videos, we refer to Deform3DGS [44] for depth estimation of the videos to obtain ground true depths of the images.

Evaluation Metrics. We evaluate under 4 standard metrics [45]–[47]: root mean squared error RMSE, mean absolute error MAE, mean absolute relative error REL in millimeters (mm), and δ accuracy, which measures the percentage of predicted pixels \hat{d} that satisfy $\max(\frac{\hat{d}}{d}, \frac{d}{\hat{d}}) < 1.25^k$, where $k = 1$. This metric provides a robust measure of the proportion of depth estimates within a clinically acceptable error margin.

B. Implementation Details

We implement the proposed method with PyTorch on 2 NVIDIA RTX 4090 GPUs for 36 epochs with a batch size of 6. The dimensions of the images used in the model are 256×320 . We employ the AdamW optimizer for optimization with the initial learning rate set to 0.001. For the depth gradient fusion, we use $i = 5$ steps. For the depth diffusion model, we use the U-Net [48] architecture as the noise estimation network, and the time step T and sampling step S for the forward diffusion and backward denoising processes were set to 1000 and 20, respectively. For fair comparison, we create sparse depth maps following [32].

C. Results

The proposed method is compared with six SOTA baselines: DepthAnything-v2 [10] and EndoDAC [22] address the task of depth estimation, focusing on conventional and surgical scenes, respectively. Meanwhile, CompletionFormer [31], Marigold-DC [42], and OGNI-DC [9] are leading methods representative of the depth completion task within conventional environments. Results are shown in Table I, and Fig.4. From the comparisons in Table I, our method outperforms all metrics on both datasets. For the accuracy metric

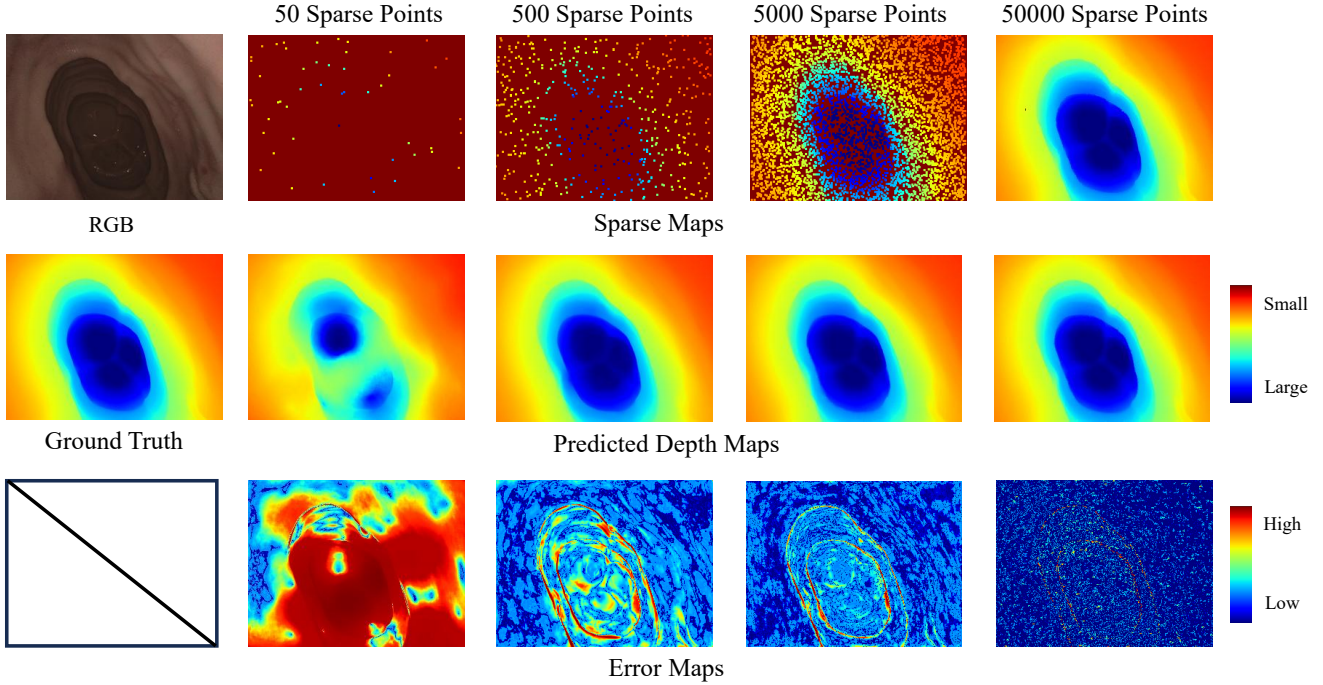


Fig. 5. Results of EndoDDC at different levels of sparsity. Our method generates robust dense depth across different levels of sparsity.

TABLE II
QUANTITATIVE COMPARISON RESULTS OF DEPTH COMPLETION MODELS AT DIFFERENT SPARSITY LEVELS ON C3VD [40] DATASET.

Methods	50		500		5000		50000	
	RMSE (mm) ↓	REL (mm) ↓	RMSE (mm) ↓	REL (mm) ↓	RMSE (mm) ↓	REL (mm) ↓	RMSE (mm) ↓	REL (mm) ↓
CompletionFormer [31]	4.9533	0.0717	0.6875	0.0071	0.5848	0.0064	5.7541	0.1629
Marigold-DC [42]	1.8955	0.0302	0.8294	0.0106	0.5596	0.0091	0.4569	0.0088
OGNI-DC [9]	3.6632	0.0767	0.6770	0.0067	0.3757	0.0029	0.1403	0.0007
Ours	3.6772	0.0703	0.6591	0.0064	0.3442	0.0025	0.1177	0.0004

TABLE III
ABLATION STUDY ON C3VD [40] AND STEREOVIS [41] DATASET. "BASELINE" INDICATES USING [31], [38] WITHOUT THE DEPTH GRAD FUSION AND DEPTH DIFFUSION. "W/O" DENOTES "WITHOUT". DDI INDICATES DEPTH-DEPTH ITERATION [9].

Methods	C3VD [40]				StereoMIS [41]			
	RMSE (mm) ↓	MAE (mm) ↓	REL (mm) ↓	δ ↑	RMSE (mm) ↓	MAE (mm) ↓	REL (mm) ↓	δ ↑
Baseline	0.6883	0.2329	0.0069	0.9987	1.6159	0.6163	0.0065	0.9985
Baseline+DDI [9]	0.6770	0.2283	0.0067	0.9988	1.5857	0.6114	0.0063	0.9986
Ours (w/o Depth Grad)	0.6669	0.2304	0.0068	0.9988	1.5729	0.5893	0.0064	0.9985
Ours (w/o Init Depth)	0.6453	0.2261	0.0063	0.9983	1.5604	0.5675	0.0062	0.9985
Ours Full	0.6412	0.2104	0.0060	0.9990	1.4691	0.5515	0.0061	0.9988

δ , our method improves compared to all other methods, especially compared to EndoDAC in the StereoMIS dataset, where our method improves the accuracy by 25.55%. For the error metrics, compared to the fine-tuned DepthAnything-v2 and self-supervised learning depth estimate models EndoDAC, our method reduces RMSE, MAE, and REL metrics substantially on both datasets, even compared to the best depth completion model OGNI-DC, RMSE, MAE, and REL on C3VD are reduced respectively by 5.28%, 7.84%, and 10.44%, while on StereoMIS, they decreased by 7.35%, 9.79%, and 3.17%, respectively.

In Fig.4, we observe that the fine-tuning DepthAnything-v2 and EndoDAC perform poorly in terms of millimeter-level depth estimation when compared to other depth completion methods. This is because, compared to the depth completion model, this class of methods lacks the constraints of a sparse depth map during training and can only estimate the depth value through the relationship between image frames, resulting in a weaker depth prediction capability compared to depth completion models. Additionally, when compared to the other three depth completion models, Marigold-DC, CompletionFormer, and OGNI-DC, our method demonstrates

superior depth prediction capabilities, especially at image edges and regions with significant depth variations. In the error maps, we can find that our method exhibits smaller estimation errors at the edges of the image and in regions with large depth gradient variations. In general, our approach outperforms the other SOTA methods.

D. Robustness to Different Sparse Maps

For downstream applications in surgical robotics, a model’s robustness to varying levels of sparsity is of critical importance [1]. However, retraining separate models for different sparsity levels is impractical for real-world applications, as the degree of depth sparsity cannot be known a priori in an actual surgical environment. For example, in robot-assisted minimally invasive surgery, the acquired sparse depth observations can fluctuate drastically due to occlusions by surgical instruments, specular reflections from tissue surfaces, or changing endoscopic viewpoints. Therefore, we aim to build a single model that performs well across all sparsity levels.

Following [32], we uniformly trained the models on the C3VD [40] under 500 sparse points and evaluated their performance at 50, 500, 5000, and 50000 sparse points, comparing them with multiple baseline models. Fig. 5 shows our model’s test results across different sparsity levels. As Table II indicates, except at 50 sparse points, our model consistently outperforms other baseline models under identical sparsity conditions. This is because EndoDDC, as an end-to-end model specifically designed for endoscopic robots, has an architecture optimized for deeply integrating RGB features with sparse geometric input information. When sparse inputs are drastically reduced to 50 points, the geometric constraints provided become extremely weak. In contrast, the baseline model of Marigold-DC – Stable Diffusion [49], trained on billions of images, leverages a powerful visual prior pre-trained on massive datasets to generate relatively well-structured depth maps even under such severely underconstrained conditions. However, the nature of the depth completion task relies not on generation driven by visual prior, but on high-precision reconstruction based on dense geometric constraints. Consequently, as the number of sparse points increases, EndoDDC demonstrates superior performance compared to other models.

E. Ablation Study

In the ablation study, we compare our method with the baseline [31], [38] and the DDI [9] method for iterative update. For fair comparison, we report the results of the baseline without the Depth Diffusion and with DD. To enable fair comparison, we fine-tune the baseline method on the surgical datasets [40], [41]. Specifically, we initialized the baseline models with weights pre-trained on natural image datasets and subsequently trained them on C3VD and StereoMIS using the same optimizer and learning rate as EndoDDC. This process ensures that the performance gap is not due to domain shift between natural and surgical

environments, but rather stems from the architectural advantages of our proposed diffusion-based refinement. We also compare our full method with the version using depth diffusion without the Depth Grad Fusion feature as diffusion condition, and the version of depth diffusion with depth gradient condition but without the initial depth. The former performs diffusion without guidance of depth and gradient features, while the latter does not incorporate initial depth information during denoising, instead initiating the diffusion process directly from random noise. The results are presented in Table III. Compared with others, our method shows dramatic improvement across all metrics. This result shows that our proposed Depth Grad Fusion module (for diffusion conditioning) and Depth Diffusion module with initial depth significantly optimize the depth map, leading to state-of-the-art performance by effectively modeling the data distribution and reducing residual errors.

V. CONCLUSIONS

In this paper, we propose EndoDDC, an advanced endoscopy depth completion method that significantly improves depth estimation accuracy by introducing sparse depth information and using depth and depth gradient features for conditioning diffusion optimization. Experimental results on two publicly available endoscopy datasets show that the proposed method outperforms the state-of-the-art in all fine-tuning methods, self-supervised depth estimation, and depth completion tasks, particularly in handling complex endoscopic environments with weak textures and challenging lighting conditions. By providing an accurate and robust depth completion method, this work has the potential to significantly enhance the autonomous navigation capabilities of surgical robots, improving procedural safety and efficiency.

REFERENCES

- [1] Z. Min, J. Lai, and H. Ren, “Innovating robot-assisted surgery through large vision models,” *Nat. Rev. Electr. Eng.*, vol. 2, pp. 350–363, 2025.
- [2] H. Saeidi, J. D. Opfermann, M. Kam, S. Wei, S. Léonard, M. H. Hsieh, J. U. Kang, and A. Krieger, “Autonomous robotic laparoscopic surgery for intestinal anastomosis,” *Sci. Robot.*, vol. 7, no. 62, p. eabj2908, 2022.
- [3] C. Budd and T. Vercauteren, “Transferring relative monocular depth to surgical vision with temporal consistency,” in *Proc. 2024 Int. Conf. Med. Image Comput. Comput.-Assist. Interv. (MICCAI)*, 2024, pp. 692–702.
- [4] D. Stellinga, D. B. Phillips, S. P. Mekhail, A. Selyem, S. Turtaev, T. Čížmár, and M. J. Padgett, “Time-of-flight 3d imaging through multimode optical fibers,” *Science*, vol. 374, no. 6573, pp. 1395–1399, 2021.
- [5] Y. Jiang, Z. Dong, and S. Mai, “Robust cost volume generation method for dense stereo matching in endoscopic scenarios,” *Sensors*, vol. 23, no. 7, p. 3427, 2023.
- [6] Z. Yang, J. Pan, J. Dai, Z. Sun, and Y. Xiao, “Self-supervised lightweight depth estimation in endoscopy combining cnn and transformer,” *IEEE Trans. Med. Imaging*, vol. 43, no. 5, pp. 1934–1944, 2024.
- [7] X. Anadón, J. Rodríguez-Puigvert, and J. Montiel, “3d densification for multi-map monocular vslam in endoscopy,” *arXiv preprint arXiv:2503.14346*, 2025.
- [8] F. Ma and S. Karaman, “Sparse-to-dense: Depth prediction from sparse depth samples and a single image,” in *Proc. 2018 IEEE Int. Conf. Robot. Autom. (ICRA)*, 2018, pp. 4796–4803.
- [9] Y. Zuo and J. Deng, “Ogni-dc: Robust depth completion with optimization-guided neural iterations,” in *Proc. 2024 Eur. Conf. Comput. Vis. (ECCV)*, 2024, pp. 78–95.

- [10] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, "Depth anything v2," in *Adv. Neural Inf. Process. Syst.*, vol. 37, 2024, pp. 21 875–21 911.
- [11] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud, "Dust3r: Geometric 3d vision made easy," in *Proc. 2024 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024.
- [12] V. Leroy, Y. Cabon, and J. Revaud, "Grounding image matching in 3d with mast3r," in *Proc. 2024 Eur. Conf. Comput. Vis. (ECCV)*, 2024, pp. 71–91.
- [13] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, "Vggt: Visual geometry grounded transformer," in *Proc. 2025 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2025, pp. 5294–5306.
- [14] K. Niu, Z. Liu, X. Feng, H. Li, Q. Lin, and K. Shi, "Endoscopic depth estimation based on deep learning: A survey," *arXiv preprint arXiv:2507.20881*, 2025.
- [15] Q. Tian, Z. Chen, H. Liao, X. Huang, L. Li, S. Ourselin, and H. Liu, "EndoOmni: Zero-shot cross-dataset depth estimation in endoscopy by robust self-learning from noisy labels," *arXiv preprint arXiv:2409.05442*, 2024.
- [16] D. Recasens, J. Lamarca, J. M. Fàcil, J. M. Montiel, and J. Civera, "Endo-depth-and-motion: Reconstruction and tracking in endoscopic videos using depth networks and photometric constraints," *IEEE Robot. Autom. Lett.*, vol. 6, no. 4, pp. 7225–7232, 2021.
- [17] J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Nr-slam: Nonrigid monocular slam," *IEEE Trans. Robot.*, vol. 40, pp. 4252–4264, 2024.
- [18] K. Wang, C. Yang, Y. Wang, S. Li, Y. Wang, Q. Dou, X. Yang, and W. Shen, "Endogslam: Real-time dense reconstruction and tracking in endoscopic surgeries using gaussian splatting," in *Proc. 2024 Int. Conf. Med. Image Comput. Comput.-Assist. Interv. (MICCAI)*, 2024, pp. 219–229.
- [19] Y. Huang, B. Cui, L. Bai, Z. Chen, J. Wu, Z. Li, H. Liu, and H. Ren, "Advancing dense endoscopic reconstruction with gaussian splatting-driven surface normal-aware tracking and mapping," in *Proc. 2025 IEEE Int. Conf. Robot. Autom. (ICRA)*, 2025, pp. 4084–4091.
- [20] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam," *IEEE Trans. Robot.*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [21] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, p. 139, 2023.
- [22] B. Cui, M. Islam, L. Bai, A. Wang, and H. Ren, "Endodac: Efficient adapting foundation model for self-supervised depth estimation from any endoscopic camera," in *Proc. 2024 Int. Conf. Med. Image Comput. Comput.-Assist. Interv. (MICCAI)*, 2024, pp. 208–218.
- [23] X. Liu, A. Sinha, M. Ishii, G. D. Hager, A. Reiter, R. H. Taylor, and M. Unberath, "Dense depth estimation in monocular endoscopy with self-supervised learning methods," *IEEE Trans. Med. Imaging*, vol. 39, no. 5, pp. 1438–1447, 2019.
- [24] J. Tang, F.-P. Tian, W. Feng, J. Li, and P. Tan, "Learning guided convolutional network for depth completion," *IEEE Trans. Image Process.*, vol. 30, pp. 1116–1129, 2020.
- [25] S. Liu, S. De Mello, J. Gu, G. Zhong, M.-H. Yang, and J. Kautz, "Learning affinity via spatial propagation networks," in *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.
- [26] X. Cheng, P. Wang, and R. Yang, "Depth estimation via affinity learned with convolutional spatial propagation network," in *Proc. 2018 Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 103–119.
- [27] X. Cheng, P. Wang, C. Guan, and R. Yang, "Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion," in *Proc. 2020 AAAI Conf. Artif. Intell. (AAAI)*, vol. 34, 2020, pp. 10 615–10 622.
- [28] J. Park, K. Joo, Z. Hu, C.-K. Liu, and I. So Kweon, "Non-local spatial propagation network for depth completion," in *Proc. 2020 Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 120–136.
- [29] Y. Lin, H. Yang, T. Cheng, W. Zhou, and Z. Yin, "Dyspn: Learning dynamic affinity for image-guided depth completion," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 6, pp. 4596–4609, 2023.
- [30] K. Rho, J. Ha, and Y. Kim, "Guideformer: Transformers for image guided depth completion," in *Proc. 2022 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2022, pp. 6250–6259.
- [31] Y. Zhang, X. Guo, M. Poggi, Z. Zhu, G. Huang, and S. Mattoccia, "Completionformer: Depth completion with convolutions and vision transformers," in *Proc. 2023 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023, pp. 18 527–18 536.
- [32] R. Liu, Z. Liu, H. Zhang, G. Zhang, Z. Zuo, and W. Sheng, "Dense depth completion based on multi-scale confidence and self-attention mechanism for intestinal endoscopy," in *Proc. 2023 IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 7476–7482.
- [33] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," *arXiv preprint arXiv:2010.02502*, 2020.
- [34] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," *arXiv preprint arXiv:2011.13456*, 2020.
- [35] T. Salimans and J. Ho, "Progressive distillation for fast sampling of diffusion models," *arXiv preprint arXiv:2202.00512*, 2022.
- [36] W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao, "Pyramid vision transformer: A versatile backbone for dense prediction without convolutions," in *Proc. 2021 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2021, pp. 568–578.
- [37] Y. Zhang, X. Guo, M. Poggi, Z. Zhu, G. Huang, and S. Mattoccia, "Completionformer: Depth completion with convolutions and vision transformers," in *Proc. 2023 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023, pp. 18 527–18 536.
- [38] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *Proc. 2020 Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 402–419.
- [39] Y. Lin, T. Cheng, Q. Zhong, W. Zhou, and H. Yang, "Dynamic spatial propagation network for depth completion," in *Proc. 2022 AAAI Conf. Artif. Intell. (AAAI)*, vol. 36, 2022, pp. 1638–1646.
- [40] T. L. Bobrow, M. Golhar, R. Vijayan, V. S. Akshintala, J. R. Garcia, and N. J. Durr, "Colonoscopy 3d video dataset with paired depth from 2d-3d registration," *Med. Image Anal.*, vol. 90, p. 102956, 2023.
- [41] M. Hayoz, C. Hahne, M. Gallardo, D. Candinas, T. Kurmann, M. Allan, and R. Sznitman, "Learning how to robustly estimate camera pose in endoscopic videos," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 18, no. 7, pp. 1185–1192, 2023.
- [42] M. Viola, K. Qu, N. Metzger, B. Ke, A. Becker, K. Schindler, and A. Obukhov, "Marigold-dc: Zero-shot monocular depth completion with guided diffusion," *arXiv preprint arXiv:2412.13389*, 2024.
- [43] J. Rodríguez-Puigvert, V. M. Batlle, J. M. M. Montiel, R. Martínez-Cantín, P. Fua, J. D. Tardós, and J. Civera, "Lightdepth: Single-view depth self-supervision from illumination decline," in *Proc. 2023 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023, pp. 21 273–21 283.
- [44] S. Yang, Q. Li, D. Shen, B. Gong, Q. Dou, and Y. Jin, "Deform3dgs: Flexible deformation for fast surgical scene reconstruction with gaussian splatting," in *Proc. 2024 Int. Conf. Med. Image Comput. Comput.-Assist. Interv. (MICCAI)*, 2024, pp. 132–142.
- [45] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 3, pp. 1623–1637, 2020.
- [46] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," in *Proc. 2021 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2021, pp. 12 179–12 188.
- [47] M. A. U. Khan, D. Nazir, A. Pagani, H. Mokayed, M. Liwicki, D. Stricker, and M. Z. Afzal, "A comprehensive survey of depth completion approaches," *Sensors*, vol. 22, no. 18, p. 6969, 2022.
- [48] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. 2015 Int. Conf. Med. Image Comput. Comput.-Assist. Interv. (MICCAI)*, 2015, pp. 234–241.
- [49] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proc. 2022 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2022, pp. 10 684–10 695.