# Bimodal Camera Pose Prediction for Endoscopy

Anita Rau, Binod Bhattarai, Lourdes Agapito, and Danail Stoyanov

***Abstract*— Deducing the 3D structure of endoscopic scenes from images remains extremely challenging. In addition to deformation and view-dependent lighting, tubular structures like the colon present problems stemming from the self-occluding, repetitive anatomical structures. In this paper, we propose *SimCol*, a synthetic dataset for camera pose estimation in colonoscopy and a novel method that explicitly learns a bimodal distribution to predict the endoscope pose. Our dataset replicates real colonoscope motion and highlights drawbacks of existing methods. We publish 18k RGB images from simulated colonoscopy with corresponding depth and camera poses and make our data generation environment in Unity publicly available. We evaluate different camera pose prediction methods and demonstrate that, when trained on our data, they generalize to real colonoscopy sequences and our bimodal approach outperforms prior unimodal work. Our project and dataset can be found here: `www.github.com/anitarau/simcol`.**

***Index Terms*— 3D reconstruction, camera pose estimation, endoscopy, SLAM, surgical AI**

## I. Introduction

Reliably reconstructing the colon during colonoscopy from raw endoscopic images could improve cancer screening quality and substantially impact patient outcomes. A 3D model could help endoscopists navigate and determine surfaces of the colon wall that were not sufficiently screened for polyps. Detecting and removing all polyps can prevent colorectal cancer from developing and is a crucial predictor of survival rates [1]. In regions lacking well-trained colonoscopists, such a technology can increase the amount of cancer screening services and their quality. Yet, there is a vast gap between the ability to predict 3D structures from images of generic scenes versus endoscopic scenes. Feature-based Structure-from-Motion (SfM) or Simultaneous Mapping And Localization (SLAM) are standard approaches for reconstructing outdoor or indoor scenes offline [2] or online [3]. Such methods detect features, establish correspondences, and estimate the relative camera motion and the 3D scene mutually and accurately. However, the application of feature-based systems to colonoscopic scenes suffers significant limitations that inhibit practical use in computer-assisted interventions (CAI). CAI systems have to address self-occlusion, camera-pose-dependent lighting and shadows, reflective surfaces, and a general lack of texture [4]. Reconstructing the colon based on features thus remains (even in offline applications) an unsolved and challenging problem [5]. Learning methods do not necessarily rely on features and therefore have great potential to further 3D reconstruction during colonoscopy. Accurately predicted camera pose and scene depth can be an initial estimate within an SfM or SLAM pipeline, provide information about certainty, or, downstream, enable the direct regression of scene coordinates and camera locations in an end-to-end fashion [14], [15]. Depth prediction during colonoscopy has already been studied and can be considered robust [16]–[18]. However, predicting camera poses in this context remains challenging.

Working toward a learning-based system requires vast amounts of training data. However, labels for real colonoscopy are not readily available: depth and pose sensors cannot fit into a standard colonoscope. Therefore, existing methods have focused on surrogate data with ground truth (simulated or porcine) or self-supervised approaches. Unfortunately, existing

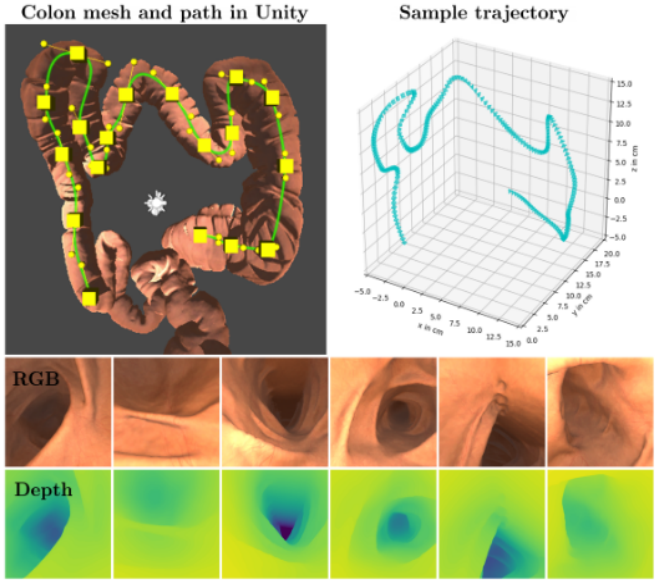


Fig. 1. Our proposed dataset provides ground truth depth maps and camera poses. We show a virtual environment; an example virtual endoscope trajectory travelling through it; and representative images with their depth maps below them, where blue represents far surfaces.

This work was supported by the Wellcome/EPSRC Centre for Interventional and Surgical Sciences (WEISS) [203145Z/16/Z]; Engineering and Physical Sciences Research Council (EPSRC) [EP/P027938/1, EP/R004080/1, EP/P012841/1]; The Royal Academy of Engineering Chair in Emerging Technologies scheme; and the EndoMapper project by Horizon 2020 FET (GA 863146). 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.

All authors are affiliated with University Collge London, Gower St, London WC1E 6BT, United Kingdom (e-mail:{a.rau.16, b.bhattarai, l.agapito, danail.stoyanov}@ucl.ac.uk).

The authors would like to thank Javier Morlana from University of Zaragoza for providing the COLMAP results for real colonoscopy sequences and both Sophia Bano from UCL and the anonymous reviewers for the constructive discussions and comments.

TABLE I
AN OVERVIEW AND COMPARISON OF COLONOSCOPY OR GASTROSCOPY DATASETS

| Dataset | Description | R/V/P | Public | Depths | Camera pose | Intrinsics | Tubular | # Frames | # Trajectories |
|---|---|---|---|---|---|---|---|---|---|
| Ozyoruk *et al.* [6] | Ex-vivo porcine 3D scanned colon | R | some | ✓ | ✓ | ✓ | × | >8k | 18 |
| Ozyoruk *et al.* [6] | Simulated capsule endoscope in *Unity* | V | some | ✓ | ✓ | ✓ | ✓ | 22k | 1 |
| Armin *et al.* [7] | Simulated colonopscoy video | V | × | ? | ✓ | ? | ✓ | 30k | >15 |
| Turan *et al.* [8] | Ex-vivo porcine stomach | R | × | × | ✓ | ✓ | × | 12k | > 4 |
| Ma *et al.* [9] | COLMAP labels from colonoscopy video | R | × | × | × | ? | ✓ | 1.2m | 60 |
| Widya *et al.* [10] | Stomach with real meshed texture | R/V | × | ✓ | ✓ | ✓ | × | ? | 7 |
| Freedman *et al.* [11] | Rendered synthetic dataset | V | × | ✓ | ? | ? | ✓ | 187k | ? |
| Fulton *et al.* [12] | Colon phantom magnetic tracker dataset | P | ✓ | × | ✓ | ✓ | ✓ | 24k | 7 |
| Bae *et al.* [13] | SfM labels from colonoscopy video | R | × | × | × | ? | ✓ | >34k | 51 |
| **Ours** | Simulated colonoscope in *Unity* | V | ✓ | ✓ | ✓ | ✓ | ✓ | 18k | 15 |

R = Real, V = Virtually simulated, P = Physical phantom, ? = Could not be verified, > could be partially verified but it appears there are more.

simulated and porcine datasets are only partially publicly available. Further, ex-vivo porcine does not look similar to in-vivo human colonoscopy, and simulated datasets do not account for the authenticity of the simulated camera movement. We argue that for pose estimation, not only must the appearance of a synthetic dataset be realistic, but the simulated camera movements also must replicate the movements of a real colonoscope during colonoscopy.

One might argue that self-supervised methods are the remedy: They do not require labels at all as they instead mutually train a depth and a pose network [6] using warping losses. However, we show that such methods are *not* guaranteed to converge to accurate depths and poses. When poses are bimodally distributed, we show that self-supervised methods can learn a distribution with one maximum, instead of two.

We, therefore, introduce a new, fully publicly available benchmark dataset that will enable the development and benchmarking of supervised camera pose prediction and 3D reconstruction methods during colonoscopy. The simulated camera poses replicate real camera movement during screening more closely than existing data and help networks generalize to real data. Additionally, we propose a novel supervised pose regression network that explicitly models a bimodal distribution through a combined classification-regression framework. The classification network classifies the camera movement into insertion or withdrawal, and the regression network predicts the difference from the respective class mean.

To work toward more accurate relative camera pose predictions during colonoscopy in this paper we:

- create and publish an extensive synthetic dataset of over 18,000 images with rendered RGB images, depth information, camera poses, camera intrinsics, and detailed documentation;
- explain why unsupervised methods are not readily applicable to infer relative camera poses accurately;
- propose a bimodal approach that more accurately predicts our dataset's bimodally distributed camera movement.

## II. RELATED WORK

The availability of real colonoscopy training data with camera poses and scene depths is limited. Existing convolutional neural nets (CNNs) thus focus mainly on synthetic data, labels from ex-vivo porcine, or unsupervised methods.

**Unsupervised depth and pose prediction:** Self-supervised methods require neither ground truth nor pseudo labels. They depend instead on carefully designed loss functions. Turan *et al.*, Ozyoruk *et al.*, Yao *et al.*, and Widya *et al.* all employ warping errors to supervise a depth and pose prediction network in gastroscopy or colonoscopy [6], [8], [19], [20]. Though an unsupervised approach enables training in an environment where it is near impossible to obtain good ground truth, we identify several shortcomings of existing approaches in this work. One drawback is that unsupervised methods can predict unimodal distributions even though the underlying data is bimodally distributed. We discuss this phenomenon in detail in Section III-A.

**Integrated CNN & SfM pipelines:** Sub-tasks of the traditional SfM pipeline can be learned with CNNs; vice-versa, SfM outputs can drive CNNs. In both scenarios, the method heavily relies on SfM pseudo labels. However, standard SfM fails more often than not in colonoscopy, and methods that rely on their outputs can only learn to be as good as these pseudo labels. Ma *et al.* [9] incorporate the output of a depth network into their SfM pipeline and predict an initial pose using a recurrent neural net that is refined within a traditional SLAM pipeline. Their CNN is trained on sparse SfM-depth. Bae *at al.* [13] propose a multi-view stereo algorithm for dense depth that is able to match patches between images but depends on an initial SfM camera pose and sparse depth estimation.

**Pose regression with CNNs:** When simulated data or ex-vivo porcine data is available, networks can learn in a supervised manner. Armin *et al.* [7] train a CNN on synthetic data to directly output the 6DoF vector describing the relative rotation and translation between two cameras during colonoscopy. The back-propagated error is a weighted sum of squared translation and rotation error. Turan *et al.* [21] incorporate information from entire video sequences of the upper GI-tract using recurrent neural networks to enable time-consistent camera pose predictions. Their network directly regresses a 6DoF vector from two consecutive RGB images and their estimated corresponding depth maps. While supervised approaches tend to be accurate, they require data that is not widely available.

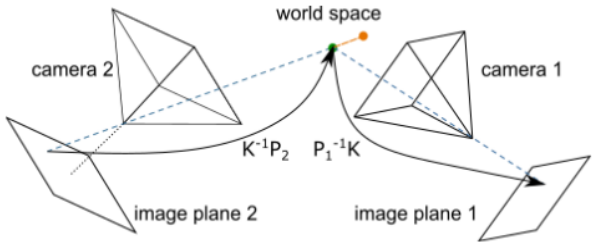


Fig. 2. Illustration of image warping. The orange point is occluded in image 2 and would not appear in warped image 1.

**Datasets for camera pose estimation in colonoscopy:** Different works have used a variety of datasets, but only a few published them. Table I enumerates the datasets referenced in relevant papers. To the best of our knowledge, the only fully available dataset for colonoscopy is based on a phantom [12] and investigates the effects of deformation. It was created using an electro-magnetic tracker attached to an endoscope that provides 6DoF poses. However, the structures are so repetitive that any forward movement could be interpreted as a backward movement looking at the previous fold rather than the next fold. Ozyoruk *et al.* created a real and a simulated dataset and made some trajectories publicly available [6]. While it is the first of its kind for real data and is immensely helpful, its usefulness to the research community could be limited. Creating the real dataset required a porcine colon to be mounted to scaffolding. The mounting prevented the visualization of the typical shape of the colon and the camera's pointing towards the lumen. As a result, the published trajectories consist mostly of colon wall images. A virtual capsule endoscope generated the simulated dataset in Unity while a user "steered" the camera through a colon mesh. While this may replicate a colonoscopy with a magnetic pill-shaped capsule, the data does not reflect the movement of a standard colonoscope. Only 0.004% of consecutive frame pairs in the published trajectory have a non-zero rotation [22]. A fully documented benchmark dataset based on the movement of a standard colonoscope with a clear distinction between training and test data is still missing.

## III. Methods

Our motivation is based on a prominent prior work on self-supervised depth and pose prediction [6]. We develop and make available a novel dataset and describe the data generation process and properties of our data. Then we propose a novel, bimodal approach to camera pose estimation that is better suited to learning bimodally distributed data. Our model, the full dataset, the documentation, and the data generation environment will be publicly available.

### *A. A study on the limitations of self-supervised depth and pose prediction in colonoscopy.*

Ideally, a network can learn reliable depth and pose estimates from image pairs. Depth and pose are inherently related through the geometry of the scene. Self-supervised approaches are based on the idea that predicted depth and relative pose between two cameras could be used to warp an image to a nearby frame. A loss function then compares the warped image to the target image. We define the pose of a camera in a reference world space at time $\tau$ as a $4 \times 4$ matrix

$$P_\tau = \begin{bmatrix} \mathbf{R}_\tau & \mathbf{t}_\tau \\ \mathbf{0} & 1 \end{bmatrix},$$

where $\mathbf{R}$ is a $3 \times 3$ rotation matrix and $\mathbf{t}$ is a $3 \times 1$ translation vector. We can project a pixel into the respective camera frame using the camera intrinsics.

As illustrated in Figure 2, we can *warp* pixels between images. We project a pixel to its camera space using the intrinsic matrix $K$. The camera's pose in the world space can then project a point $h = [x, y, z, 1]$ in homogeneous coordinates into world space. The inverse of a second pose lets us project the point into the new camera space from where $K$ maps it to the new image plane. We summarize:

$$h' = P_1^{-1} P_2 h =: \Omega h, \tag{1}$$

where $h'$ is the image of a point $h$ in a new camera frame and $\Omega$ is a $4 \times 4$ projection matrix. It is important to note that world points occluded by structures in one image can not be retrieved through warping. For a thorough discussion of projective geometry, we refer to Hartley and Zisserman [23].

Similarly to projecting points between camera spaces, we can map a camera pose to a new pose in the world space. In other words, we rotate and translate a camera. To map camera 1 to the pose of camera 2, we write

$$P_2 = (P_1 P_1^{-1}) P_2 = P_1 (P_1^{-1} P_2) = P_1 \Omega. \tag{2}$$

In the context of mapping camera poses we refer to $\Omega$ as the *relative camera pose*. Self-supervised methods technically learn the warping parameter $\Omega$, but it can be interpreted as the relative camera pose.

To understand how self-supervised methods learn $\Omega$, we define the relevant losses that compare a target image to the predicted warped frame. Let Images $I_\tau$, $I_{\tau+k}$ be two nearby RGB images that are $k$ frames apart. For synthetic data, let $D_\tau$, $D_{\tau+k}$ be the corresponding ground truth depth maps and let $\Omega_{\tau \to \tau+k}$ be the projection matrix that projects points from camera frame $\tau$ to camera frame $\tau+k$. Further, let hats denote the predictions of a network: $\hat{D}_\tau$ and $\hat{\Omega}_{\tau \to \tau+k}$. Note that networks can parameterize camera translations and rotations differently, but other representations can be mapped to $4 \times 4$ projection matrices [24].

Then image $I_\tau$ can be warped to look like $I_{\tau+k}$ and the warped image is denoted through a tilde as $\tilde{I}_{\tau+k}(\Omega_{\tau+k \to \tau}, D_{\tau+k}, I_\tau)$, based on the inverse projection, target depth, and reference image used. For dense warping, we use *inverse warping*; that is, we warp the target image back to the reference image rather than warping the reference to the target. The reprojection loss for an image pair $(I_\tau, I_{\tau+k})$ is defined as

$$L_R = \sum_{i=-k,k} \sum_\tau ||I_{\tau+i} - \tilde{I}_{\tau+i}(\hat{\Omega}_{\tau+i \to \tau}, \hat{D}_{\tau+i}, I_\tau)||_1. \tag{3}$$

Other relevant losses are the geometric consistency loss $L_G$ [6] that compares the warped depth $\tilde{D}_{\tau+k}$ to the projected depth

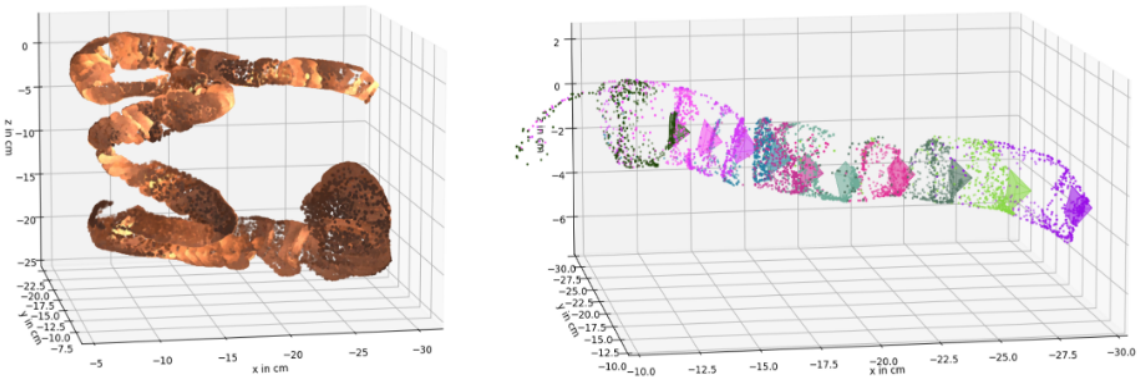


Fig. 3. Projecting ground truth depth and camera poses into world space replicates the colon mesh. Left: Every tenth frame of one test trajectory is projected into the same world space. Right: Every twentieth frame of a smaller section in the same test sequence is projected into world space together with the respective camera poses. Each camera and its point cloud have the same random color, and clouds are subsampled for better visualization.

$D'_{\tau+k}$, and the structural similarity index measure (SSIM), here referred to as $L_S$, as defined in [25].

But due to the geometry of the colon, and the related characteristic camera movement, mapping between image plane and 3D world is ambiguous (and certainly not bijective):

**Illumination inconsistency:** Camera-pose-dependent lighting interferes with warping losses: The same surface can look very different from slightly different positions. Adjusting for illumination inconsistency by scaling the intensity values of an image pair to have the same mean [6] only linearly approximates the illumination differences. It neglects, for instance, that colon folds throw shadows and block light from reaching farther parts of the lumen. Illumination is an important cue for understanding geometry in the mostly textureless colon. Its inconsistency challenges the applicability of warping-based methods in the colon.

**Smoothness inertia:** Networks tend to predict smooth depth maps when trained with naive warping loss, and this is often even reinforced by a smoothness loss [26]. While this might be a reasonable approach for smooth surfaces like bronchi, the geometry of the colon is characterized through folds and thus sharp steps in depth.

**Back-projection ambiguity:** Due to self-occlusion warping errors and geometric inconsistency often are not minimized by the ground truth depth and pose—the network converges to a local minimum.

**Wide field of view:** The wide frustum of colonoscopes emphasizes nearby structures over-proportionally in the image.

We show experimentally how these drawbacks can lead to a sub-optimal accuracy of the predicted camera pose.

### B. SimCol: A new dataset

We propose a new public dataset that will enable pose and depth prediction in colonoscopy. Following the pipeline introduced in [17] we use a *Unity* simulation environment; however, we also generate ground truth camera poses. Our data is based on a computer tomography scan of a human colon from which a textured mesh was created in [6].

We define a path in the center of the colon mesh defined through 18 `WayPoints` which the virtual colonoscope follows, simulating a traversal through a colon. Our dataset aims to reflect typical camera movements during colonoscopy. During withdrawal in a real-world procedure, the scope is slowly pulled backwards along the colon's centerline. When folds present themselves, the tip is flexed and rotated to observe the entire mucosa. Reinsertion and re-withdrawal often occur [27]. Similarly, the virtual camera follows the centerline of the colon while rotating around it. The camera outputs images with size 475x475 and replicates the intrinsics of a real colonoscope. Two virtual light sources are attached to the left and the right of the camera. Each time a user renders images along the trajectory, the `WayPoints` are independently offset by a random translation of up to 2mm, and a rotation of up to 20 degrees, which we chose experimentally.

When generating new data, a user can change the rotation, position, randomization, and number of Waypoints arbitrarily. The user can also adjust the interpolation parameters between Waypoints. Adjusting a Waypoint will affect all frames' camera poses between the previous and the following Waypoint. The camera intrinsics, image resolution, and light sources can also be adjusted in our Unity project. Given all poses as 4x4 projection matrices describing the camera's pose in the world reference frame, we assume the first camera ($\tau$) of each pair to be at the origin. Thus we learn the projection $\Omega_\tau = P_\tau^{-1} P_{\tau+k}$, that describes the position of camera 2's ($\tau + k$) origin as seen from camera 1. Accordingly, a positive z-coordinate of $\Omega$ indicates that camera 2 is in front of camera 1. As *Unity* uses a left-handed system, while the *Python* package *Scipy* uses a right-handed system, we transform the *Unity* camera poses during training and testing using the projection $P^{right} = MP^{left}M$, where $M$ is a $4 \times 4$ identity matrix with $-1$ as the second diagonal element. After transformation, the z-axis points forward from the optical center, while the y-axis points upward and the x-axis to the right. Figure 1 shows example images, depths, and trajectories next to the colon mesh in its *Unity* environment. The Figure shows the camera path in green and the `WayPoints` in yellow. Algorithm 1 gives an overview of the data generation process. Furthermore, we show how the depth maps and poses can be used to project RGB images into 3D space in Figure 3(a), and we visualize the ground truth camera poses of a smaller section and their respective

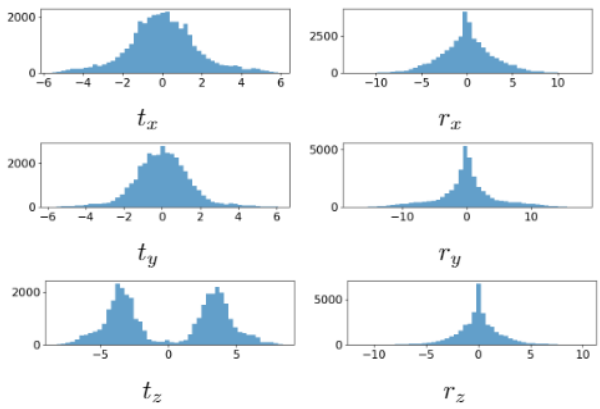


Fig. 4. Histograms of relative $x$-, $y$-, and $z$- translations (t) in millimeters and rotations (r) in degrees between image pairs in our dataset that are $k = 5$ frames apart. The translation along the z-axis is clearly bimodal. This observation corresponds to the endoscope's predominant forward and backward movement within the colon.

**Algorithm 1** Data generation in Unity

```
for wayPoint in WayPoints do
    randomize wayPoint pose
    time = 0
    while time < timePerWayPoint do
        cam.position ← GetPosition(wayPoint, time)
        cam.rotation ← GetRotation(wayPoint, time)
        render RGB image
        render depth map
        update time
    end while
end for
```

point clouds in (b). Lastly, we plot histograms of the x-, y-, and z-rotations and translations for consecutive frame pairs in Figure 4. It is important to note that the translation along the z-axis is clearly bimodal, as we would expect in a real scenario. Further, the rotations range in an interval of about [-2,2] degrees. Our dataset is therefore the first public dataset with verified ground truth depths and poses, both training and testing data, and a wide variety in rotation.

### C. A new approach to predict bimodal camera pose

Inspired by [28] we propose the use of an explicitly bimodal model to learn bimodally distributed camera pose.

Our model uses a ResNet-18 architecture as encoder and, for comparability with [6], a 4-layer convolutional pose decoder. The bimodal approach is visualized in Figure 5 and explained in the following.

**Class net:** We propose a classification network that classifies relative poses into insertion $t_z > 0$ or withdrawal $t_z < 0$ movements, where $t_z$ denotes the translation along the $z$ axis based on the projection matrix $\Omega_\tau = P_\tau^{-1} P_{\tau+k}$. In other words, the translation along the $z$-axis is classified into two *bins*. Our network first embeds each image of a pair $(I_\tau, I_{\tau+k})$ into a $256 \times 16 \times 16$ feature vector using a ResNet-18 architecture referred to as $f$. The embedded images are then concatenated along the feature dimension and passed into both the classification net $\mathcal{C}$ and a regression net $\mathcal{R}$. The class net

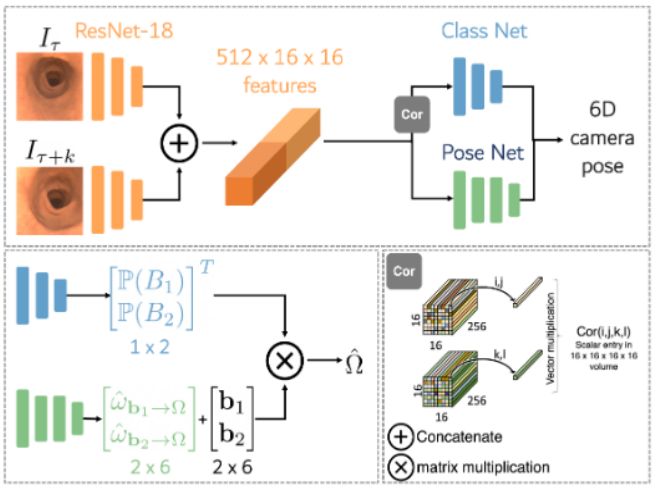


Fig. 5. Architecture of our bimodal model. Top: Overview of the proposed approach. Lower left: Outputs of the Pose Net and Class Net constitute the predicted relative pose $\hat{\Omega}$. lower right: Illustration of the correlation layer.

is a simple 3-layer fully concatenated network with dropout layers with high dropout probability (0.5) as proposed in [28]. It outputs a 2D vector that is passed through a softmax layer and can be interpreted as the probability $\mathbb{P}$ that the input image is from bin $B_1$ or bin $B_2$, respectively:

$$\begin{aligned}\mathcal{C}(f(I_\tau), f(I_{\tau+k})) &= [\mathbb{P}(t_z < 0) \quad \mathbb{P}(t_z > 0)] \\ &= [\mathbb{P}(\Omega_\tau \in B_1) \quad \mathbb{P}(\Omega_\tau \in B_2)],\end{aligned}$$

where *B1* and *B2* represent two sets of relative poses $P$, such that $P \in B_1 \iff P^{-1} \in B_2$. As we assume that each class is normally distributed we train the classification network with cross entropy loss

$$L_{class} = -\sum_i \mathbb{P}(B_i) \log \hat{\mathbb{P}}(B_i). \tag{4}$$

**Representation of relative pose:** There are different ways to represent camera poses. As a rotation matrix has many almost zero entries, it is not a good representation for a network to learn. Therefore, we use the logarithms of unit quaternions in this work. They are represented as a 3D vector that can be mapped to the 4D unit quaternion as proposed in [29]. The pose error

$$L_{pose} = |\hat{t} - t| \exp^{-\beta} + \beta + |\log \hat{q} - \log q| \exp^{-\gamma} + \gamma \tag{5}$$

is a weighted sum of rotation error and translation error, where the weights $\beta$ and $\gamma$ are learned as proposed in [30] and adapted to log quaternions in [31] with initial values 0 and $-3$, respectively. The final pose output of our network is a 6D vector representing the 3D translation and the 3D log quaternion. For convenience, we continue to refer to the relative pose in the new representation as $\Omega$.

**Regression net:** The regression net $\mathcal{R}$ is a simple 4-layer fully convolutional layer with Relu activations as used in [6]. $\mathcal{R}$ outputs a 6D pose estimate for each of the two classes relative to each of the two bins as proposed in [28]. Let *bin1*, *bin2* be the approximate centers of $B_1$ and $B_2$ describing the

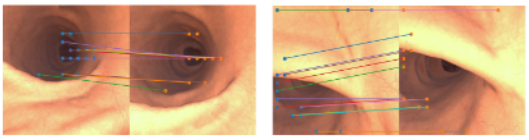

Fig. 6. Visualization of the correlation layer: The argmax let's us illustrate the pixels with the highest correlation in the other image.

translation along the z-axis and $\mathbf{b}_1 = [0, 0, bin1, 0, 0, 0]^T$. Then $\mathcal{R}$ predicts the difference between the ground truth pose in 6D representation and both $\mathbf{b}_1$ and $\mathbf{b}_2$. We denote the differences as $\hat{\omega}_{\mathbf{b}_1 \to \Omega}$ and $\hat{\omega}_{\mathbf{b}_2 \to \Omega}$. The regression net $\mathcal{R}$ outputs

$$\mathcal{R}(f(I_\tau), f(I_{\tau+k})) = [\hat{\omega}_{\mathbf{b}_1 \to \Omega_\tau} \quad \hat{\omega}_{\mathbf{b}_2 \to \Omega_\tau}]. \tag{6}$$

**Final prediction:** The final pose prediction is a weighted sum of the probabilities from the class net and the regressed pose differences

$$\hat{\Omega}_\tau = \begin{bmatrix} \mathbb{P}(\Omega_\tau \in B_1) \\ \mathbb{P}(\Omega_\tau \in B_2) \end{bmatrix}^T \cdot \begin{bmatrix} \mathbf{b}_1 + \hat{\omega}_{\mathbf{b}_1 \to \Omega_\tau} \\ \mathbf{b}_2 + \hat{\omega}_{\mathbf{b}_2 \to \Omega_\tau} \end{bmatrix}. \tag{7}$$

**Correlation layer:** Crucial to the performance of the class net is the correlation layer. As introduced in [32] and used for pose regression in [33], the down-sampled images are passed through a layer in which the correlation between each feature pair is computed. The output is a $16 \times 16 \times 16 \times 16$ correlation volume. The volume is flattened and normalized along the first two dimensions to indicate matches of $I_{\tau+k}$ in $I_\tau$, and along the last two dimensions for the opposite case. Unlike [32], [33] our correlation layer has no learnable weights. Yet, it is crucial for the convergence of the classification net as it decouples the features for pose regression from the class net: The class net does not depend on the learned features but only on the similarity across images. The correlation layer matches features between an image pair, and although they are too noisy to be useful for direct pose estimation, their signal helps distinguish between forward and backward movements easily. In fact, the class net has reached an accuracy of 99% on the validation set after the first epoch. For visualization only, we return the argmax along the feature dimension and receive a 16x16 grid with the location of the most similar feature in the other image in each pixel. We upsample the grid matches to the centers of their respective patches in the $475 \times 475$ image and show the 15 matches with the highest correlation in Figure 6. Experimenting with the argmax as input to the class net, we found the volume more useful.

**Target function:** Let $w_c$ be the weight for the class loss. Then our network is trained minimizing the target function

$$L = L_{pose} + w_c L_{class}. \tag{8}$$

## IV. Experiments

In this section, we first observe the drawbacks of self-supervised pose networks. We then show that models trained on our data can better generalize to real colonoscopy sequences. Moreover, we show that supervised methods, too, can generalize to real data, making them suitable to be used for real procedures and lending themselves to domain adaptation methods. Lastly, we show that our novel bimodal method outperforms the unimodal baseline and is better suited to learning the camera movement during real colonoscopy.

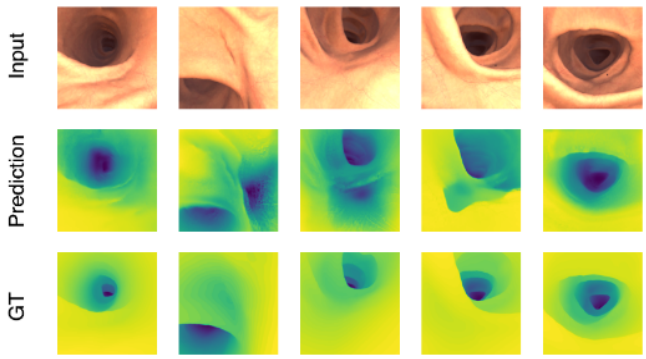


Fig. 7. Examples for predicted and ground truth (GT) depth maps using the self-supervised baseline. Scales are omitted as depths are predicted up to scale. Blue corresponds to a higher depth.

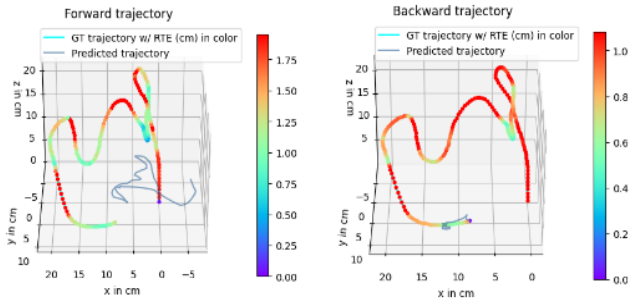


Fig. 8. Predicted camera trajectories using the self-supervised baseline are worse for the backward traversal than for the forward traversal. RTE is color coded in the ground truth trajectory.

### A. Minima of self-supervised target functions are not guaranteed to minimize depth and pose errors

We train the method proposed and published in [6] on our own data with all hyper-parameters set to their default values. To obtain the full trajectory of absolute poses, we multiply the relative poses. The pose of a camera $\tau$ in world space can be computed as $P_\tau = P_1 \Omega_1 \cdot \ldots \cdot \Omega_{\tau-1}$, where each $\Omega$ projects the initial pose $P_1$ sequentially to the subsequent camera pose. As the network predicts pose up to scale we measure the accuracy of the *scaled* trajectory using the Absolute Translation Error (ATE), the Relative Translation Error (RTE), and the ROTation error (ROT). The losses are defined as

$$RTE = \mu_\tau(||trans({\Omega_\tau}^{-1} \hat{\Omega}_\tau)||), \tag{9}$$

$$ATE = \mu_\tau(||trans(P_\tau) - trans(\hat{P}_\tau)||), \quad \text{and} \tag{10}$$

$$ROT = \mu_\tau(\frac{trace(Rot({\Omega_\tau}^{-1} \hat{\Omega}_\tau)) - 1}{2} \cdot \frac{180}{\pi}), \tag{11}$$

where $\mu_\tau$ denotes the median over all steps $\tau$, $trans$ and $Rot$ denote the translation and rotation components of a projection matrix, and $||.||$ denotes the 2-norm. The ATE measures drift and the overall consistency of the predicted trajectory; however, it is prone to outliers. More robust than the ATE is the RTE. It measures the magnitude of the difference between

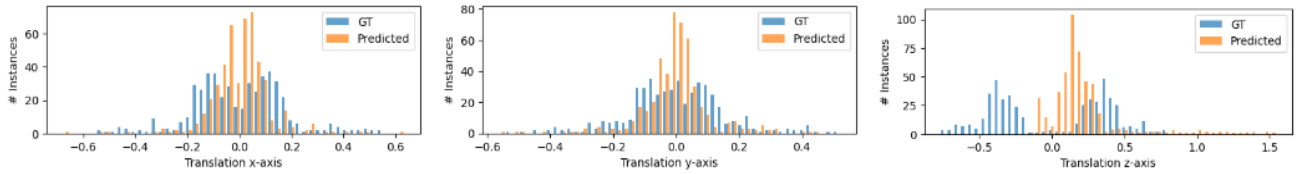


Fig. 9. Minimizing standard self-supervised losses is not sufficient to learn a bimodal distribution. Histograms of predicted vs. ground truth (GT) poses on test set 1. Translations in cm.

TABLE II
ABSOLUTE TRANSLATION ERRORS, RELATIVE TRANSLATION ERRORS, AND ROTATION ERRORS ON OUR TEST TRAJECTORIES

| Test trajectory | 1 | | | | 2 | | | | 3 | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Total trajectory length | 105.1 cm | | | | 103.6 cm | | | | 102.2 cm | | | |
| Mean step size | 4.4 mm | | | | 4.3 mm | | | | 4.3 mm | | | |
| Mean rotation per step | 4.6° | | | | 4.8° | | | | 4.8° | | | |
| | ATE(cm) | RTE(mm) | ROT(°) | Acc(%) | ATE(cm) | RTE(mm) | ROT(°) | Acc(%) | ATE(cm) | RTE(mm) | ROT(°) | Acc(%) |
| COLMAP [2] | 0.01 | 0.07 | 0.18 | 8 | 0.02 | 0.10 | 0.28 | 6 | 0.37 | 0.68 | 1.42 | 7 |
| EndoSLAM [6] | 14.9/17.8 | 3.10/4.26 | 1.3/5.7 | 100/41 | 18.0/15.3 | 3.50/4.28 | 1.8/5.0 | 100/35 | 14.9/15.7 | 3.24/4.16 | 1.6/5.5 | 100/37 |
| Ours unimodal | 7.08/12.5 | 0.75/0.79 | **1.3/1.3** | 100/99 | **2.79/2.63** | 0.76/0.76 | **1.5/1.5** | 100/100 | 7.56/10.0 | 0.87/0.90 | 1.5/1.6 | 100/99 |
| Ours bimodal | 8.86/13.3 | 0.72/0.72 | 1.5/1.5 | 100/99 | 2.83/6.03 | 0.69/0.71 | 1.7/1.7 | 100/100 | 6.17/9.62 | 0.85/0.89 | 1.6/1.6 | 99/99 |
| **Ours bimodal** w/$L_c$ | **8.81/9.79** | **0.69/0.72** | 1.5/1.5 | 100/99 | 2.35/5.23 | **0.67/0.70** | 1.6/1.6 | 100/100 | **3.78/9.00** | **0.82/0.85** | 1.6/1.5 | 100/99 |
| Ours bimodal w/$L_c$ w/o Corr | 14.0/13.9 | 0.77/0.76 | 1.7/1.6 | 99/99 | 8.32/5.46 | 0.73/0.71 | 2.1/2.0 | 99/98 | 8.70/8.10 | 0.93/0.98 | 1.9/1.8 | 99/98 |

ATE & RTE, ROTation error, and accuracy for our test trajectories' forward/backward traversal. COLMAP performs global optimization and thus direction does not matter (only one value reported). The accuracy for regression methods (Ours and EndoSLAM) denotes the percentage of correctly predicted directions. The accuracy for COLMAP reports the percentage of frames from the trajectory that COLMAP was able to reconstruct. Note that although COLMAP yields the smallest errors, the method successfully reconstructs only a small fraction of each trajectory. Bold indicates the best result as the sum of forward and backward errors. W/o Corr refers to our model without the correlation layer in the classfication net.

the predicted and actual relative pose per step and reflects translation and rotation errors on a local level. As we do not apply global optimization or loop-closure, and all methods in this work predict each relative pose independently, the RTE has more significance for us. Lastly, the ROT measures only the magnitude of the rotation of the local errors.

The scaling factor is defined as

$$s = \frac{\sum_\tau trans(P_\tau)^T \cdot trans(\hat{P}_\tau)}{\sum_\tau trans(\hat{P}_\tau)^T \cdot trans(\hat{P}_\tau)}, \tag{12}$$

where $P_\tau$ denotes absolute camera poses. A different approach would be to compute an alignment projection for the prediction, as in [34]. However, as we investigate the ability to predict correct directions, we only scale the predicted trajectory but refrain from translating or rotating it. If a network were to predict the exact opposite of the ground truth, a rotation of 180° could compensate for the error, which is not informative in our case. We compute independent scales for the forward and backward trajectories.

In Figure 7, we show that the network has learned to predict smooth yet sensible depth maps. Because the network is trained based on the self-supervision of warped RGB and depth images, one might assume that the network must have learned both depth and camera pose. Investigating the resulting trajectory as shown in Figure 8, we find that the network predicts the insertion movement considerably better than the opposite movement. The forward movement successfully predicts the three half-loops, though with drift. Predicting the motion in the opposite direction does not lead to a sensible trajectory. Further evidence that the performance of the algorithm varies depending on direction is presented in Figure 9. The histogram of predicted z-values is not aligned with the ground truth values leading to a bias towards a positive translation along the z-axis. Finally, also Table II (EndoSLAM method) underlines our findings, where we quantitatively evaluate the performance of the self-supervised approach on three different trajectories within our dataset. Most notably, the RTE and rotation errors are higher for the withdrawal than for the insertion for all trajectories.

**Dissecting insertion vs. withdrawal:** We use the terms insert and withdraw when speaking of direction relative to the structure rather than forward/backward because the latter can be defined arbitrarily. So what is the difference between inserting and withdrawal movements? During training, the network predicts the forward and backward movement between an image pair $(I_\tau, I_{\tau+k})$. We show warped images in both directions based on ground truth depth and pose vs. warped images based on predicted (ˆ) depth and pose in Figure 10. We also show error maps and the resulting training losses. Observe that column (viii) is minimized during training to look like column (vii). Moreover, column (viii) should look like column (ix) if the predicted depth and pose were correct. Note that even when using the ground truth values, the training losses are non-zero. Visually, there are significant differences between the ground truth warp and the predicted warp, especially along steps in the structure. Investigating the resulting errors in column (xii) shows high errors for geometry consistency losses (D).

For the opposite movement, the images in columns (ii) and (iii) should look like column (i). We observe that the ground truth warped images have many artifacts that originate in the self-occlusion of the scene and the sharp differences in depth. These artifacts lead to high warping errors for the ground truth,

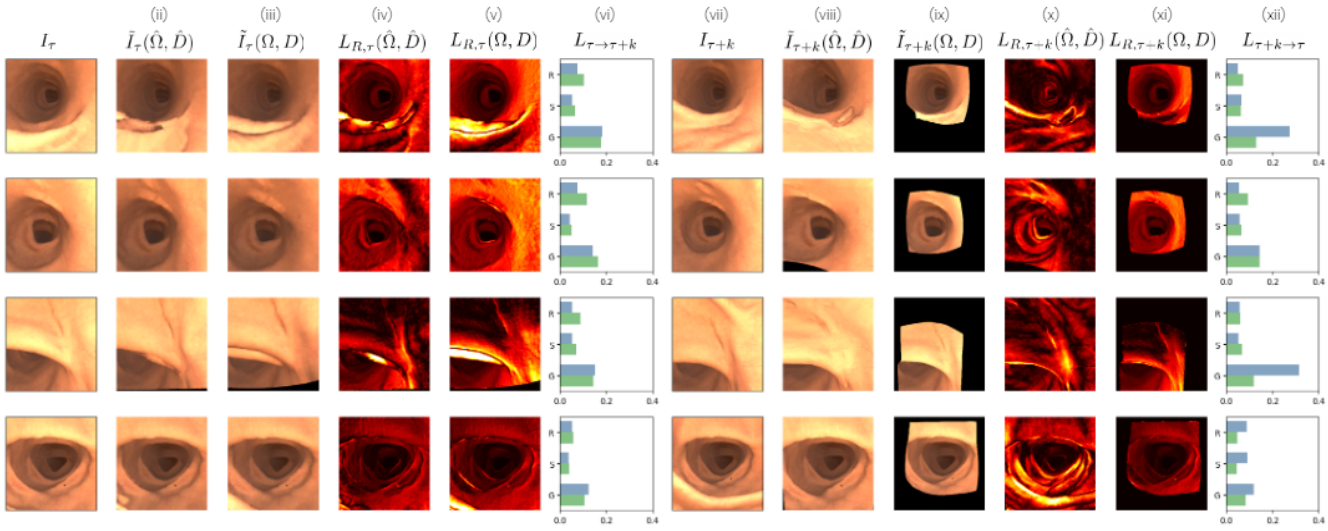


R=Reprojection loss, S=SSIM loss, G=geometric consistency loss.

Fig. 10. Training losses of a self-supervised depth and pose network are not necessarily minimized at ground truth. Each row shows an image pair $(I_\tau, I_{\tau+k})$ and its induced training losses. The first six columns visualize the errors of warping $I_{\tau+k}$ to $I_\tau$, while the last six rows show the opposite case: (i) + (vii) original images; (ii), (iii), (viii), (ix) warped images $\tilde{I}_{\tau+k\to\tau}$ and $\tilde{I}_{\tau\to\tau+k}$ using predicted depth and camera pose vs. using ground truth depth and pose; (iv), (v), (x), (xi) error maps for the Reprojection loss $L_R$, where dark areas represent low errors; and (vi), (xii) $L_R, L_S$ and $L_G$ for the predicted and ground truth inputs. If depth and pose were correctly learned we would expect (ii) to equal (iii), and (viii) to equal (ix). However, especially (viii) fails to represent the correct camera movement. Further, we would expect that using ground truth depth and camera pose leads to lower warping errors in (vi) and (xii) than using the predictions; yet, the errors induced by the ground truth are often higher (green bars larger than corresponding blue bars). In all examples the camera moves in negative z-direction between frames $\tau$ and $\tau + k$.

as shown in columns (vi). The predicted warpings visually replicate nearby structures well, resulting in minor reprojection errors, although wrong poses are predicted. Every example in column (vi) has a higher reprojection error for the ground truth (green) than for the prediction (blue). The network has converged to a wrong optimum. We found that the network predicts higher reprojection losses for the ground truth than for the predictions in 63% of the test images. We observed a similar, though weaker, behavior using warp errors based on HSV color model or depth. However, we found the RGB-based reprojection loss to be crucial for the convergence.

We conclude that, due to the geometry of the colon, depth prediction and pose prediction do not necessarily improve or help each other. The different properties of the insertion and the withdrawal warping lead to different performances for the pose prediction. As a result, the unsupervised network learns a unimodal distribution of the z-translation shifted towards positive z-translations. The model in this case was not able to predict both forward and backward movements of the camera. Because these are the main relative movements to be expected during colonoscopy, the bimodality needs to be accounted for.

### B. Our data helps generalize to real data

In this section, we show that methods trained on our *SimCol* data better generalize to real trajectories. We evaluate the results of three methods on six real colonoscopy sequences. We hand-select sequences that were successfully reconstructed with COLMAP [2]. Upon visual examination, we found most reconstructed sequences to produce flat or noisy point clouds indicating errors in the reconstruction. We picked sequences that reflected a reasonable and coherent 3D structure of the observed surface. Because we rely on visual inspection, the COLMAP results are not considered ground truths; however, they are a useful baseline.

TABLE III
ROBUSTNESS AND GENERLIZABILITY TO UNSEEN REAL DATA

| | (a) | (b) | (c) | (d) | (e) | (f) |
|---|---|---|---|---|---|---|
| Ttl* | 12.0 | 14.6 | 11.6 | 11.9 | 11.1 | 15.0 |
| Mss* | 0.7 | 0.4 | 0.9 | 0.7 | 0.6 | 0.8 |
| Mrs | 2.1 | 1.5 | 2.1 | 1.6 | 2.3 | 2.0 |
| ATE → | 1.2 3.5 7.9 | 1.4 6.2 8.4 | 1.3 1.8 11. | 1.6 1.6 5.6 | 2.6 2.3 3.2 | 1.5 6.5 7.5 |
| ATE ← | 1.0 3.2 2.3 | 1.2 1.8 3.6 | 0.7 1.2 2.1 | 1.1 6.6 6.5 | 1.1 11. 9.6 | 1.1 6.5 6.3 |
| RTE → | 0.3 0.4 1.0 | 0.2 0.3 0.6 | 0.2 0.3 1.7 | 0.2 0.3 0.5 | 0.3 0.4 0.5 | 0.2 0.5 0.7 |
| RTE ← | 0.3 0.4 0.5 | 0.2 0.1 0.2 | 0.3 0.3 0.6 | 0.3 0.5 1.1 | 0.2 1.0 1.0 | 0.3 0.4 0.5 |
| ROT → | 2.0 1.7 2.0 | 1.2 1.0 1.8 | 3.2 1.9 2.4 | 1.0 1.5 1.4 | 2.9 2.4 2.0 | 1.1 1.8 1.2 |
| ROT ← | 2.0 0.9 2.7 | 1.1 1.0 1.4 | 3.1 1.9 2.4 | 1.0 2.8 1.9 | 3.2 2.5 2.6 | 1.3 1.5 1.1 |

Ttl = Total trajectory length, Mss = Mean step size, Mrs = Mean rotation per step in degrees. Best method underlined.

Comparison of three methods evaluated on six real sequence: (i) Our proposed method trained on our proposed dataset, (ii) EndoSLAM method trained on our dataset *SimCol*, and (iii) EndoSLAM method trained on EndoSLAM data . ATE & RTE in unknown scale and Rotation error in degrees for the forward (→) and backward (←) traversal. Errors relative to COLMAP labels. * Absolute scale unknown.

In Table III, we compare three different algorithms: (i) our proposed bimodal method trained on our new dataset, (ii) the EndoSLAM algorithm using the publicly available network weights that were trained on EndoSLAM data, and (iii) the EndoSLAM algorithm trained on our data. We show the predicted poses in Figure 11. We scale trajectories resulting from the poses according to Equation 12. The errors in Table III describe the difference between COLMAP results and scaled predicted results.

Let us first compare the EndoSLAM (ES) method trained

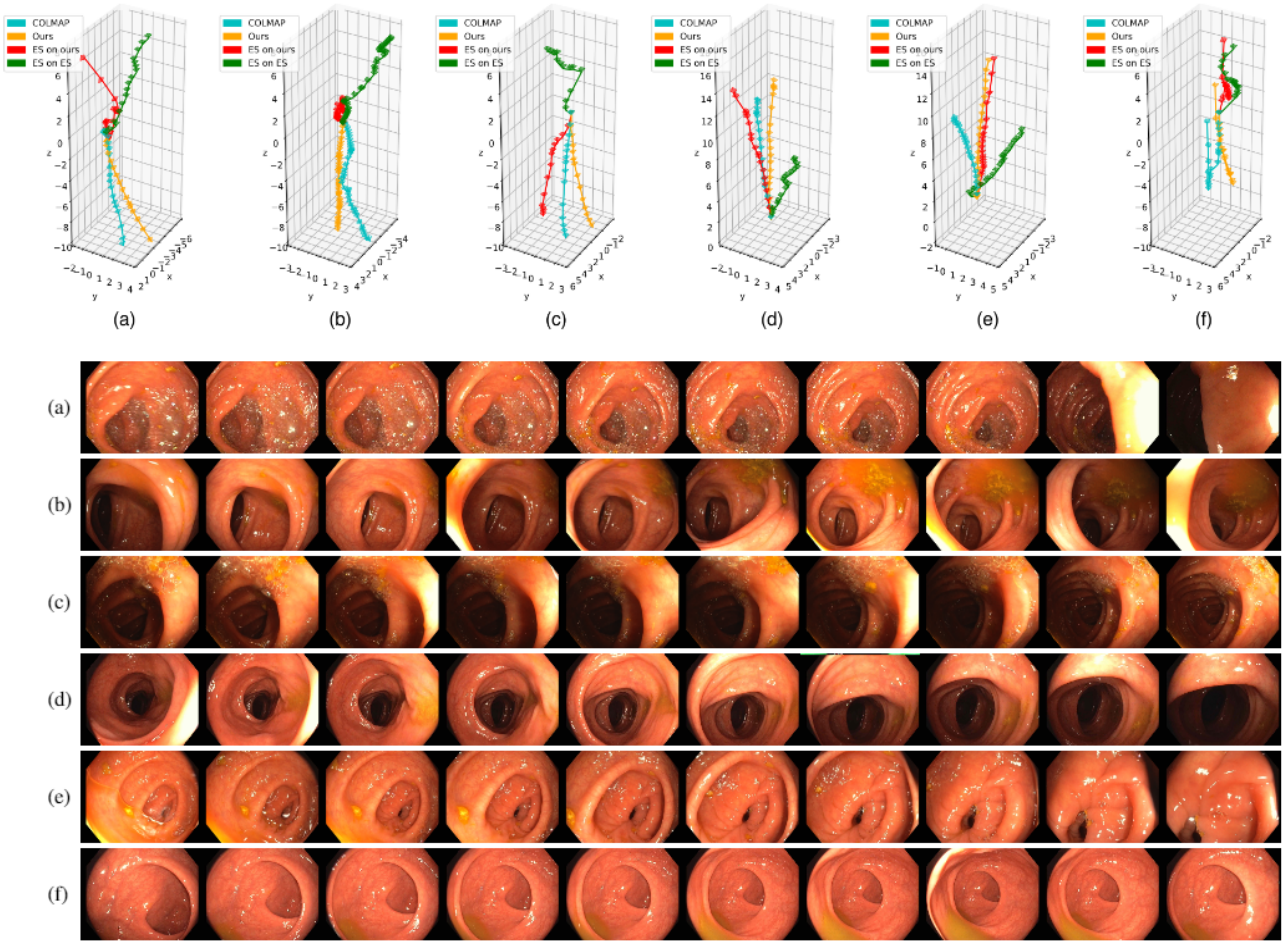


Fig. 11. Comparison of different methods with COLMAP on three withdrawal sequences (a-c), two insertion sequences (d-e), and one withdrawal followed by insertion (f). Top: Predicted and COLMAP trajectories. Bottom: Ten images from each sequence. The self-supervised methods (green and red) fail to predict most withdrawal trajectories and predict insertions instead. If trained on our data, the self-supervised method correctly predicts withdrawal sequence (c). Our method (orange) robustly predicts the direction along the z-axis in all cases, even when the camera changes direction as in (f).

on two different datasets; we will discuss the performance of our novel bimodal approach in the next section. In Table III, we can observe that ES trained on our data yields a smaller RTE in eleven out of twelve cases compared with the same method trained on the ES dataset. More importantly, in three cases, ES trained on ES fails when the same method trained on our method does not. On sequences (a), (b), and (c), the RTE produced by ES on ES is larger than the respective mean step size (Mss). That means that the error per step is larger than the step itself on average. ES trained on our data better generalizes to the camera movement in the real sequence than the same method trained on the ES data. Although the domains are different, our data provided the network with more realistic camera poses to learn. Nonetheless, both ES methods fail to consistently predict the camera movements in both the original and the reverse direction. In the withdrawal sequences (a)-(c), RTE $\rightarrow$ is especially large. Inversely, in the insertion sequences (d) - (e), RTE $\leftarrow$ is particularly large. We plot the corresponding trajectories in the original direction in Figure 11. The first three sequences show withdrawals. While ES trained on ES fails to predict the correct direction in all three cases, ES trained on our data correctly interprets the movements as a withdrawal in sequence (c). Sequences (d) and (e) show insertions. All methods suffer from drift, which we expect as we predict poses sequentially. However, ES trained on ES produces sharp direction changes in (d) when COLMAP's reconstruction is smooth. The last sequence shows a withdrawal followed by an insertion. Again, both ES baselines fail to predict the correct direction and the change in direction.

At first glance, it might be surprising that a synthetic dataset helps generalize to real data. Indeed, there is a domain gap between the appearance of real images and our synthetic data. However, our data has closed the gap between camera movements in existing datasets and camera movements in real colonoscopy videos. Our data allows a model to learn a correct prior for relative poses that is robust to the appearance gap.

### *C. Bimodal distributions are more accurately learned with bimodal models*

We previously observed how EndoSLAM performs if trained on two different datasets. Let us now observe how our

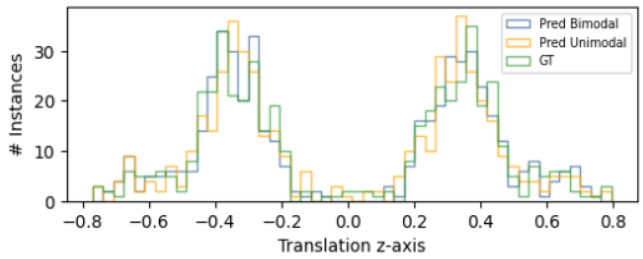


Fig. 12. Histograms of z-translations in cm for our bimodal model, unimodal model, and ground truth (GT) on test trajectory 1.

TABLE IV
MANHATTAN (L1-) ERRORS FOR Z-TRANSLATION PREDICTED BY DIFFERENT MODELS ON OUR THREE TEST TRAJECTORIES

| | 1 | 2 | 3 |
|---|---|---|---|
| Bimodal, class supervision | 132 | **122** | **154** |
| Unimodal, no supervision | 168 | 152 | 162 |
| Bimodal, no supervision | **118** | 142 | 170 |

bimodal method performs on real data. Our bimodal method trained on our dataset (orange) in Table III yields the smallest RTE and ATE in almost all cases. Unlike ES, our method yields a similar RTE for the forward and backward directions. Our model robustly predicts the trajectories in Figure 11. Especially in sequences (a), (b), and (f), where ES fails, our method accurately follows the withdrawal and even the subsequent insertion of the camera in the last case. Also, on the insertion sequences (d) and (e), the bimodal model predicts accurate trajectories.

Let us next evaluate our method on our proposed bimodally distributed synthetic data. In contrast to real data, our dataset provides labels for training and testing, which would make it unfair to compare a supervised method to a self-supervised one. We, therefore, compare our method to COLMAP. We apply COLMAP with standard settings to our data and report its performance in Table II (COLMAP). COLMAP only reconstructs a small subsection of images (14-18 images per trajectory corresponding to 6-8% of frames). Accordingly, the ATE reflects only a small subsection, and we should not compare it to the ATE of other methods. The RTE and rotation error on reconstructed subsections of test sets 1 and 2 are extremely small. The RTE and rotation error on test trajectory 3 is comparable to ours, though our approach is robust and predicts the entire trajectory. CNNs for pose prediction in colonoscopy can thus be considered a practical alternative where feature-based methods fail to initialize. Next, we evaluate the usefulness of a bimodal model compared to a unimodal approach in two experiments. We compare three different versions: (i) a bimodal model with class supervision; (ii) a unimodal model that passes the ResNet-output to the regression net only; and (iii) a bimodal model that is *not* trained with class labels. We implemented all models in PyTorch. For the bimodal models, we use the hyper-parameters $bin1 = 0.1 \cdot k$, $bin2 = -0.1 \cdot k$, $w_c = 0.1$, and $k = 5$, where $w_c$, and $k$ were chosen based on grid search.

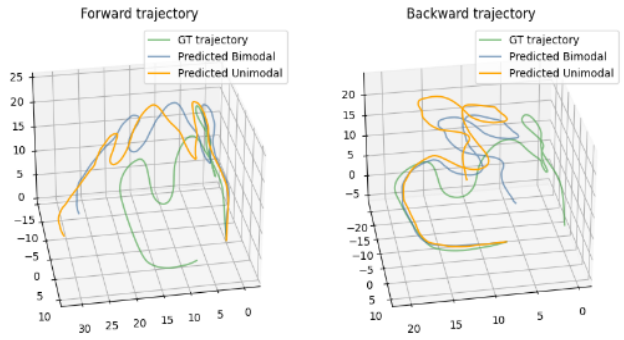


Fig. 13. Comparison of predicted trajectories produced by unimodal and our bimodal model and ground truth (GT) on test set 1. Please see supplementary material for videos.

The unimodal and bimodal models have the same number of weights in the regression net. To compare the three methods, we compare the resulting histograms and pose errors.:

Firstly, we plot histograms of the predicted poses in Figure 12. Table IV evaluates the similarity of the ground truth and predicted histograms using the Manhattan (L1-) loss

$$L_{Manhattan} = \sum_i |h(i)_{gt} - h(i)_{pred}|. \quad (13)$$

From Table IV we gather that the bimodal model supervised with cross-entropy more accurately learns the distribution of z-translations in the test set than a unimodal model for all three test trajectories. All three models, however, closely replicate the true distribution (see Figure 12).

Secondly, we evaluate the RTE, ATE, and ROT as defined in Equations 10-11 on three test trajectories in Table II. However, we do not rescale the results according to Equation 12, as we expect the network to learn the absolute scale. As we estimate the pose between a frame pair that is $k = 5$ frames apart, we calculate the errors for the five trajectories that begin with step 0, or 1, 2, 3, 4, respectively, and report the average. For each trajectory, we predict poses in forward and backward directions. The bimodal approach with supervision yields the smallest RTE for each test trajectory and, therefore, the most accurate local camera poses. However, even the bimodal approach without class supervision yields smaller RTEs than the unimodal approach on all scenes. The network learns to distinguish bins without supervision, speaking to the ability of the correlation layer to distill information about relative movements and its potential applicability to unsupervised methods. The higher error for our bimodal approach with class supervision but without correlation layer further supports the claim. Lastly, RTEs are similarly accurate in both forward and backward directions. To visualize the ATE, we plot whole trajectories for our bimodal model and the unimodal counterpart in Figure 13. The total length of the trajectory is roughly 1 meter, so drift can and should be expected.

### D. Limitations

We propose a dataset that replicates the movement of a colonoscope during real colonoscopy and observe that methods

trained on this data more robustly generalize to real video. Our bimodal approach further improves supervised methods that regress relative camera pose directly. Yet, the applicability to real colonoscopy is still limited. For a full 3D reconstruction, not only accurate poses and depths are required, but also global optimization and loop closure. Drift is universal when using local methods to predict global structures. Computational limitations are yet to be overcome to allow an on-the-fly reconstruction to give immediate feedback to colonoscopists for a higher clinical impact.

## V. Conclusions

Relative camera pose prediction during colonoscopy remains exceptionally challenging. This work explains the limitations of self-supervised methods based on warping losses for camera pose prediction. We propose using supervised methods trained on synthetic data and show that the explicit modeling of a bimodal model, rather than a unimodal model, can improve the accuracy of relative camera pose methods. Supervised methods have the additional benefit of being more robust than standard feature-based SfM methods. With our published dataset, we close the domain gap between real and synthetic relative camera poses. Future work will solve the domain gap between real and synthetic colonoscopy images, for instance, with generative adversarial networks. An exciting direction for further work could also be the investigation of bimodality within a self-supervised framework.

## References


[1] M. F. Kaminski *et al.*, "Quality indicators for colonoscopy and the risk of interval cancer," *New England Journal of Medicine*, vol. 362, no. 19, pp. 1795–1803, 2010.

[2] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[3] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "Orb-slam: a versatile and accurate monocular slam system," *IEEE transactions on robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.

[4] N. Mahmoud, A. Hostettler, T. Collins, L. Soler, C. Doignon, and J. Montiel, "Slam based quasi dense reconstruction for minimally invasive surgery scenes," *arXiv preprint arXiv:1705.09107*, 2017.

[5] O. L. Barbed, F. Chadebecq, J. Morlana, J. Martínez-Montiel, and A. C. Murillo, "Superpoint features in endoscopy," *arXiv preprint arXiv:2203.04302*, 2022.

[6] K. B. Ozyoruk *et al.*, "Endoslam dataset and an unsupervised monocular visual odometry and depth estimation approach for endoscopic videos," *Medical image analysis*, vol. 71, p. 102058, 2021.

[7] M. A. Armin, N. Barnes, J. Alvarez, H. Li, F. Grimpen, and O. Salvado, "Learning camera pose from optical colonoscopy frames through deep convolutional neural network (cnn)," in *Computer assisted and robotic endoscopy and clinical image-based procedures*, pp. 50–59, Springer, 2017.

[8] M. Turan *et al.*, "Unsupervised odometry and depth learning for endoscopic capsule robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1801–1807, IEEE, 2018.

[9] R. Ma, R. Wang, S. Pizer, J. Rosenman, S. K. McGill, and J.-M. Frahm, "Real-time 3d reconstruction of colonoscopic surfaces for determining missing regions," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 573–582, Springer, 2019.

[10] A. R. Widya, Y. Monno, M. Okutomi, S. Suzuki, T. Gotoda, and K. Miki, "Whole stomach 3d reconstruction and frame localization from monocular endoscope video," *IEEE journal of translational engineering in health and medicine*, vol. 7, pp. 1–10, 2019.

[11] D. Freedman *et al.*, "Detecting deficient coverage in colonoscopies," *arXiv preprint arXiv:2001.08589*, 2020.

[12] M. J. Fulton, J. M. Prendergast, E. R. DiTommaso, and M. E. Rentschler, "Comparing visual odometry systems in actively deforming simulated colon environments," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4988–4995, IEEE, 2020.

[13] G. Bae, I. Budvytis, C.-K. Yeung, and R. Cipolla, "Deep multi-view stereo for dense 3d reconstruction from monocular endoscopic video," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 774–783, Springer, 2020.

[14] C. Tang and P. Tan, "Ba-net: Dense bundle adjustment network," *arXiv preprint arXiv:1806.04807*, 2018.

[15] C. Zhao, L. Sun, P. Purkait, T. Duckett, and R. Stolkin, "Learning monocular visual odometry with dense 3d mapping from dense 3d flow," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6864–6871, IEEE, 2018.

[16] F. Mahmood, R. Chen, and N. J. Durr, "Unsupervised reverse domain adaptation for synthetic medical images via adversarial training," *IEEE transactions on medical imaging*, vol. 37, no. 12, pp. 2572–2581, 2018.

[17] A. Rau *et al.*, "Implicit domain adaptation with conditional generative adversarial networks for depth prediction in endoscopy," *International journal of computer assisted radiology and surgery*, vol. 14, no. 7, pp. 1167–1176, 2019.

[18] S. Mathew, S. Nadeem, S. Kumari, and A. Kaufman, "Augmenting colonoscopy using extended and directional cyclegan for lossy image translation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4696–4705, 2020.

[19] H. Yao, R. W. Stidham, Z. Gao, J. Gryak, and K. Najarian, "Motion-based camera localization system in colonoscopy videos," *Medical image analysis*, vol. 73, p. 102180, 2021.

[20] A. R. Widya, Y. Monno, M. Okutomi, S. Suzuki, T. Gotoda, and K. Miki, "Learning-based depth and pose estimation for monocular endoscope with loss generalization," *arXiv preprint arXiv:2107.13263*, 2021.

[21] M. Turan, Y. Almalioglu, H. Araujo, E. Konukoglu, and M. Sitti, "Deep endovo: A recurrent convolutional neural network (rcnn) based visual odometry approach for endoscopic capsule robots," *Neurocomputing*, vol. 275, pp. 1861–1870, 2018.

[22] K. B. Ozyoruk *et al.*, "Endoslam dataset," 2021. Data retrieved from Mendeley Data, `https://data.mendeley.com/datasets/cd2rtzm23r/1` on 27.10.2021.

[23] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[24] J.-L. Blanco, "A tutorial on se (3) transformation parameterizations and on-manifold optimization," *University of Malaga, Tech. Rep*, vol. 3, p. 6, 2010.

[25] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.

[26] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1851–1858, 2017.

[27] D. K. Rex, "Missed neoplasms and optimal colonoscopic withdrawal technique," in *Colonoscopy principles and practice*, pp. 339–350, Blackwell, Massachusetts, 2003.

[28] S. Mahendran, H. Ali, and R. Vidal, "A mixed classification-regression framework for 3d pose estimation from 2d images," *arXiv preprint arXiv:1805.03225*, 2018.

[29] S. Brahmbhatt, J. Gu, K. Kim, J. Hays, and J. Kautz, "Geometry-aware learning of maps for camera localization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2616–2625, 2018.

[30] A. Kendall and R. Cipolla, "Geometric loss functions for camera pose regression with deep learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5974–5983, 2017.

[31] M. O. Turkoglu, E. Brachmann, K. Schindler, G. J. Brostow, and A. Monszpart, "Visual camera re-localization using graph neural networks and relative pose supervision," in *2021 International Conference on 3D Vision (3DV)*, pp. 145–155, IEEE, 2021.

[32] I. Rocco, M. Cimpoi, R. Arandjelović, A. Torii, T. Pajdla, and J. Sivic, "Neighbourhood consensus networks," *arXiv preprint arXiv:1810.10510*, 2018.

[33] X. Ding, Y. Wang, L. Tang, Y. Jiao, and R. Xiong, "Improving the generalization of network based relative pose regression: dimension reduction as a regularizer," *arXiv preprint arXiv:2010.12796*, 2020.

[34] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 573–580, IEEE, 2012.