

ProbeMDE: Uncertainty-Guided Active Proprioception for Monocular Depth Estimation in Surgical Robotics

Britton Jordan^{*1}, Jordan Thompson^{*1}, Jesse F. d’Almeida², Hao Li³, Nithesh Kumar², Susheela Sharma Stern², James Ferguson¹, Ipek Oguz³, Robert J. Webster III², Daniel Brown¹, and Alan Kuntz^{3,4}

Abstract—Monocular depth estimation (MDE) provides a useful tool for robotic perception, but its predictions are often uncertain and inaccurate in challenging environments such as surgical scenes where textureless surfaces, specular reflections, and occlusions are common. To address this, we propose ProbeMDE, a cost-aware active sensing framework that combines RGB images with sparse proprioceptive measurements for MDE. Our approach utilizes an ensemble of MDE models to predict dense depth maps conditioned on both RGB images and a sparse set of known depth measurements obtained via proprioception, where the robot has touched the environment in a known configuration. We quantify predictive uncertainty via the ensemble’s variance and measure the gradient of the uncertainty with respect to candidate measurement locations. To prevent mode collapse while selecting maximally informative locations to propriocept (touch), we leverage Stein Variational Gradient Descent (SVGD) over this gradient map. We validate our method in both simulated and physical experiments on central airway obstruction surgical phantoms. Our results demonstrate that our approach outperforms baseline methods across standard depth estimation metrics, achieving higher accuracy while minimizing the number of required proprioceptive measurements. The project website is brittonjordan.github.io/probe_mde

I. INTRODUCTION

Depth perception is a fundamental requirement for robotic systems operating in unstructured and deformable environments such as surgical settings. Monocular depth estimation (MDE) offers one potential solution due to its ability to infer dense 3D scene structure from a single RGB camera without the need for specialized depth sensors [1]. However, MDE models often struggle with ambiguities arising from textureless surfaces, specular reflections, and occlusions, all of which are especially prevalent in robotic surgical settings [2].

One avenue to address these limitations is to combine vision-based predictions with proprioceptive measurements [3], [4]. We define robotic proprioception as a robot’s




Fig. 1: Conceptual illustration of our approach to fuse probabilistic learned depth models with touch-based position data. With no probed points (bottom, left), depth estimation can be uncertain. Our approach chooses the acquisition point that will minimize variance, and then re-estimates iteratively.

ability to sense contact with external objects in addition to an awareness of its own position in space through tool-tip kinematics. Utilizing proprioception, we can provide sparse ground-truth depth measurements at contact points between the robot and the environment. While these measurements can provide accurate estimates of depth at specified locations, their acquisition has a cost as each proprioceptive measurement requires physical interaction with the environment that will slow down task execution. Thus, we seek to minimize the number of required measurements while maximizing the reconstruction accuracy of the scene.

To solve this problem, we introduce ProbeMDE, a cost-aware active sensing framework for monocular depth estimation. Our approach leverages an ensemble of U-Net-based MDE models trained to predict dense depth maps from RGB images of a surgical scene conditioned on a sparse set of propriocepted depth measurements. We quantify predictive uncertainty by computing the variance across ensemble predictions to produce a differentiable scalar uncertainty measure. By analyzing the gradient of this measure with respect to potential depth measurement locations, we then employ Stein Variational Gradient Descent (SVGD) to optimize the selection of query points while minimizing mode collapse.

We validate our framework in both simulated and physical experiments. In simulation, where full ground-truth depth maps are available, we benchmark our approach against several baseline point selection methods including random

^{*}These authors contributed equally.

¹The Robotics Center and the Kahlert School of Computing at the University of Utah, Salt Lake City, UT 84112, USA.

²The Departments of ²Mechanical Engineering, ³Computer Science, and ⁴Electrical and Computer Engineering at Vanderbilt University, Nashville, TN, 37203, USA.

Research reported in this publication was supported by the Advanced Research Projects Agency for Health (ARPA-H) under Award Number D24AC00415-00 for ALISS. The ARPA-H award provided 90% of total costs with an award total of up to \$10,741,534.20 and the National Science Foundation (NSF) Graduate Research Fellowship Program funds 10% of total costs under Grant No. 2139322. The content is solely the responsibility of the authors and does not necessarily represent the official views of ARPA-H or the NSF.

sampling, greedy heuristics, and gradient-based heuristics. Our method consistently achieves greater reconstructive performance across standard depth estimation metrics. In physical experiments inside surgical phantoms, we demonstrate our method’s reconstructive accuracy and robustness over baseline approaches.

Our main contributions are twofold. First, we present a cost-aware active sensing framework that conditions monocular depth estimation on sparse proprioceptive depth measurements, including the first use of proprioception to improve monocular depth estimation in robotic endoscopic surgery. Second, we introduce a principled and efficient strategy for selecting depth measurement locations that leverages SVGD applied to the gradient of ensemble predictive variance, enabling targeted sensing that improves depth estimates and efficient measurement acquisition.

II. RELATED WORK

A. Monocular Depth Estimation

Monocular depth estimation uses an image from a single camera viewpoint to predict the distance from the camera to each point in the 3D scene. State-of-the-art MDE methods produce depth maps with outstanding accuracy for a wide variety of everyday indoor and outdoor scenes [5], [6]. However, endoscopic videos of surgical scenes are out of distribution from the large-scale datasets these foundation models are trained on, resulting in suboptimal performance [7].

Efforts to improve MDE in endoscopy include adapting or fine-tuning foundation models with simulation data [7], [8], [9], using a style transfer network to make simulation renderings more photorealistic [10], or utilizing point cloud data from structure from motion as sparse-ground-truth depth cues [11]. MDE models focused on a specific domain frequently use an adaptation of the U-Net architecture [12], [10]. In this work, we draw from a similar toolbox of approaches while adding the ability to adaptively improve depth prediction through uncertainty-guided active sensing. As in prior works, we utilize a style-transfer network to make simulation RGB renderings more realistic. We then use this data to fine-tune a large foundation model, Depth Anything V2 [5], which we subsequently use to generate pseudo-ground-truth data for training a model with better real-world performance. We introduce a novel source of cues for endoscopic MDE, robotic proprioception at specific locations within the scene which will maximally resolve the uncertainty of our prediction.

B. Uncertainty Quantification

As deep neural networks become more complex and less interpretable, uncertainty quantification has become a critical component of any such system, especially for safety-critical tasks such as robotic surgery. Many methods exist for quantifying predictive uncertainty in deep learning including deep ensembles [13], Monte Carlo dropout [14], [15], [16], and test-time augmentation [17]. A comparison of these methods shows that deep ensembles consistently provide the most accurate quantification of uncertainty under distribution

shift in the data [18]. An ensemble is created by training multiple models with different random weight initializations such that they each converge to a different local minima in the loss space. At inference time, we can use the variance of the ensemble members’ predictions as a proxy for uncertainty [19], [20], [21]. Many works have quantified the uncertainty of MDE predictions [22], [23], [24], including some focused specifically on endoscopic images [25]. However, no past works have explored the active refinement of MDE predictions using robotic touch sensing. In this work, we leverage uncertainty quantification techniques for ensembles to develop an uncertainty-guided active touch sensing framework for monocular depth estimation.

C. Active Touch Sensing

Active touch sensing has been widely explored as a means of acquiring task-relevant information under sensing constraints. Early approaches focused on reconstructing object geometry through sparse and targeted contact interactions. For instance, tactile fingertips have been used to actively trace object contours for shape recognition [26], while force-based probing has enabled reconstruction of unknown objects [27]. More recent work has emphasized uncertainty-driven exploration, where sensing actions are guided to reduce ambiguity [28]. Recent research has further advanced active touch sensing with machine learning and task-specific exploration. Learning-based approaches have been proposed to determine not only how to explore unknown objects, but also learning where to touch and how to recognize objects from tactile data [29], [30]. Together, these works demonstrate the value of active touch sensing to improve perception and task performance. In contrast, our work extends the concept of active touch sensing into the domain of proprioceptive depth acquisition, where sparse depth measurements are strategically chosen to reduce uncertainty in monocular depth estimation for deformable surgical environments.

III. PROBLEM DEFINITION

We formalize the task of selecting a sparse set of proprioceptive measurements as a cost-constrained uncertainty minimization problem for monocular depth estimation. Let $I \in \mathbb{R}^{H \times W \times 3}$ be the RGB image of a surgical scene, and let $S = \{(x_i, y_i, d_i)\}_{i=1}^N$ be a sparse set of ground-truth depth measurements where each location (x_i, y_i) corresponds to a pixel location in I and $d_i \in \mathbb{R}$ is the propriocepted depth measurement at that location. For each $s_i \in S$, there exists an associated acquisition cost $c_i > 0$. Let $D_\theta(I, S)$ be a Monocular Depth Estimator (MDE) that predicts a full depth map of the scene.

Given an uncertainty function $U(D_\theta(I, S))$ that quantifies the predictive uncertainty of the MDE, our objective is to find S that minimizes the total acquisition cost while simultaneously minimizing uncertainty. This can be expressed as the minimization of the following constrained optimization problem:

$$\min_{S \subset \Omega} \sum_{i=1}^N c_i \quad \text{s.t.} \quad U(D_\theta(I, S)) \leq \epsilon \quad (1)$$

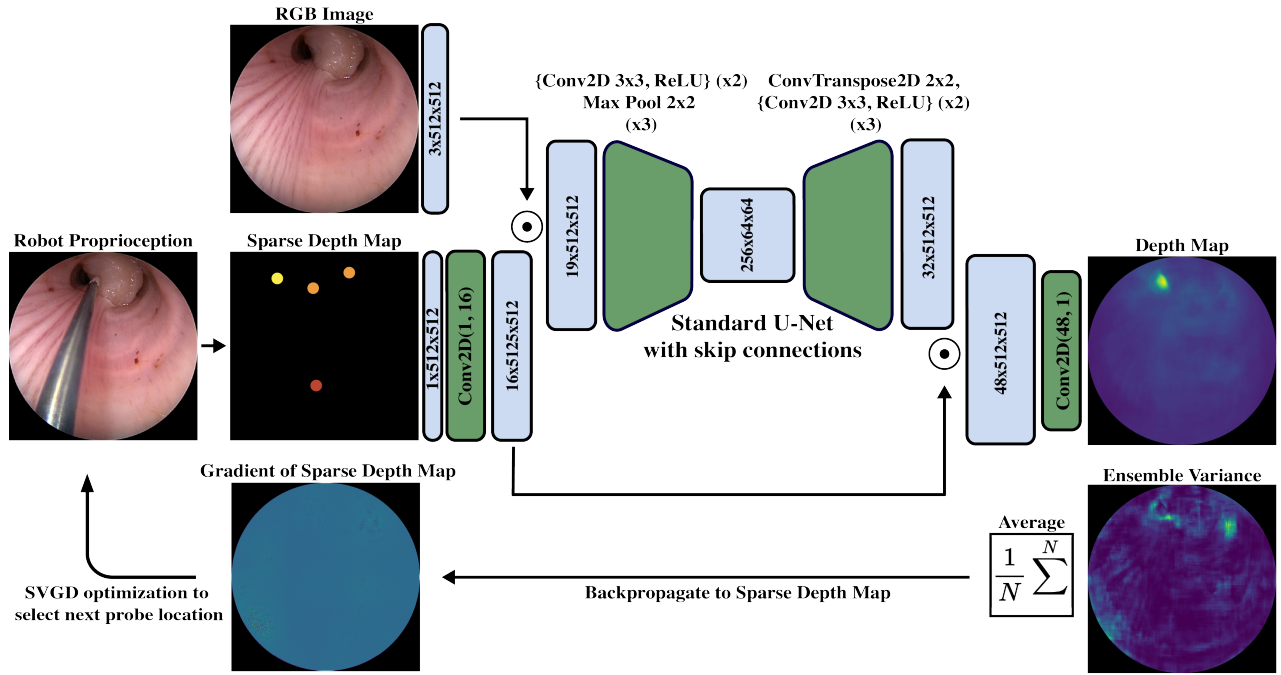


Fig. 2: ProbeMDE, a cost-aware active sensing framework for monocular depth estimation. Inputs are an endoscopic RGB image and a sparse propriocepted depth map S . We upsample S to 16 channels with a 2D convolution before concatenating with the image and passing through a U-Net encoder. Using 3 successive iterations of 2D convolutions, ReLU activation, and max pooling, the input is transformed to a 64×64 matrix with 256 channels. The convolution transpose is used to expand back to 512×512 while a 2D convolution and ReLU activation are used to combine skip connections. The upsampled sparse ground truth depth is concatenated to the output of the U-Net’s last layer before a final convolution, giving the network additional opportunity to leverage this ground-truth information. The network produces a dense 512×512 depth map. We compute the variance map U of the ensemble’s output, from which we compute $\nabla_S \|U\|$. Locations to probe next are selected using SVGD on the gradient map and then propriocepted with the surgical robot. The new depths are added to the sparse ground truth and the process is re-run.

where ε is a target threshold for the uncertainty of the MDE and Ω represents the power set of all candidate pixel locations.

IV. METHODOLOGY

A. Monocular Depth Estimation

Our MDE is based on a U-Net architecture which takes as input the RGB image of a surgical scene as well as a sparse set of ground-truth depth measurements [12]. The sparse measurements are represented as a depth map where unobserved pixels are set to -1, enabling the network to condition its predictions on available propriocepted depths. The MDE then predicts a full depth map of the scene, leveraging both visual features in the RGB image as well as the sparse depth information.

We train an ensemble of MDEs on simulation data where full ground-truth depth information is available. During training, sparse ground-truth sets are randomly sampled from the full depth map at varying densities, simulating the availability of partial propriocepted measurements. This encourages the network to incorporate known depth information when available while still learning to make accurate predictions from features in the RGB image. By training an ensemble with different initializations, we can measure the variance across predictions as a measure of model uncertainty which can subsequently be used to gather the most informative depth information.

B. Uncertainty Quantification

To quantify the uncertainty of the MDE ensemble, we first compute a pixelwise uncertainty map over the ensemble’s set of predictions. Let $\{D_\theta^{(k)}(I, S)\}_{k=1}^K$ denote the outputs of K independently trained MDE models, each conditioned on the RGB image I and the sparse ground-truth set S . The predictive variance at each pixel (x, y) is

$$U_{x,y} = \frac{1}{K} \sum_{k=1}^K (D_\theta^{(k)}(I, S)_{x,y} - \bar{D}_{x,y})^2 \quad (2)$$

where \bar{D} is the mean of the set of predictions,

$$\bar{D} = \frac{1}{K} \sum_{k=1}^K D_\theta^{(k)}(I, S). \quad (3)$$

To obtain a scalar measure of overall uncertainty, we take the mean of the full variance map

$$U_{\text{total}} = \frac{1}{N} \sum_{x,y} U_{x,y} \quad (4)$$

where N is the total number of pixels in the image. This scalar represents the total predictive uncertainty of the ensemble over the entire predicted depth map.

Since U_{total} is differentiable with respect to the sparse depth measurements in S , we can compute $\nabla_S U_{\text{total}}$ which quantifies the expected reduction in ensemble uncertainty for each potential propriocepted depth measurement. We

backpropagate through the network to the sparse depth map input, calculating the gradient of the variance average with respect to each pixel in the map. Locations corresponding to the largest-magnitude negative gradients indicate where acquiring depth information would most effectively reduce uncertainty. This gradient map therefore forms the foundation of our cost-aware active depth sensing framework, guiding the selection of depth measurements to maximize the improvement in predictive depth while minimizing the cost of gathering depth information.

C. Active Depth Sensing

After computing the gradient of the total ensemble uncertainty with respect to the sparse depth measurements, $\nabla_S U_{\text{total}}$, we use this information to actively select which locations to propriocept. Rather than selecting points greedily, we adopt Stein Variational Gradient Descent (SVGD) to optimize a set of candidate locations [31]. Selecting points greedily could lead to a set of query points all near each other. SVGD strikes a balance between clustering points in highly informative regions while also encouraging diversity throughout the image. We first negate the gradient map so that the maximum values in the map correspond with locations with the highest likelihood of reducing uncertainty. The resulting map is then normalized to form a discrete probability distribution over all candidate pixel locations, indicating where measurements would be the most informative.

Formally, let $\{z_i\}_{i=1}^N$ denote a set of N particles representing potential locations to propriocept. At each iteration, each particle is updated according to

$$z_i \leftarrow z_i + \eta \phi(z_i), \quad (5)$$

where η is the step size and $\phi(z_i)$ is the SVGD update function:

$$\phi(z_i) = \frac{1}{N} \sum_{j=1}^N \left[k(z_j, z_i) (\nabla_{z_j} \log(p(z_j)) + \nabla_{z_j} k(z_j, z_i)) \right]. \quad (6)$$

Here, $p(z_i)$ is the discrete probability distribution obtained from the normalized negative gradient map, and $k(\cdot, \cdot)$ is the Radial Basis Function (RBF) kernel to encourage spatial diversity of the particles. This prevents all of the particles from collapsing to the same location, ensuring that the propriocepted locations are diverse throughout the scene.

Through these iterative updates, the particles converge to locations that maximize the expected reduction in total MDE uncertainty. The selected points are then propriocepted and added to the sparse ground-truth set S . A new prediction can then be made using this updated input to yield a full depth map that is more accurate in previously uncertain regions, while maintaining cost efficiency in the number of measurements acquired.

V. EXPERIMENTS

A. Datasets

Simulated Dataset: A simulated dataset is generated from CT scans of 13 real-world ex vivo porcine tracheal models with an intraluminal tumor surrogate (chicken tissue) affixed




Fig. 3: Physical experiment setup with Virtuoso Endoscopy System motor pack at center and capacitance sensing module at bottom left. A probe wire runs through the concentric tubes and extends from the tip of the robot. Detail view shows the endoscope sheath inserted into an ex vivo pig trachea.

using cyanoacrylate adhesive. We convert the scans into 3D meshes in Unity. A simulated camera explores the interior of the trachea reconstructions, collecting pairs of RGB images and depth maps. The RGB images are then translated into the style of real phantom endoscopic scenes using a style-transfer network [32] to reduce the domain gap.

Real-world Dataset: Our real-world dataset is collected from 12 other ex vivo porcine tracheal models with tumor surrogates prepared as above. We collect exploratory endoscopic videos of these models using the Virtuoso Endoscopy System (Virtuoso, Nashville, TN, USA) [33] and then generate pseudo-ground-truth depth maps using Depth Anything V2’s indoor large model [5] fine-tuned on the simulated dataset. We create additional data points by linearly and exponentially stretching the values of the pseudo-ground-truth depth maps. This artificially exaggerates the ambiguity of perceived scale, thereby forcing the model to rely heavily on the sparse ground truth. These pseudo-labels are used only to supervise training of the real-world model and are not used for evaluation.

B. Ensemble Training

We train an ensemble of $K = 3$ MDE models based on a U-Net backbone. Each model is trained independently with different random initializations on the same training and testing dataset. Training is performed using both the simulated and real-world datasets producing two ensembles (one for the simulated data experiments and one for physical trials). To train on sparse ground-truth sets, for each point in

the dataset, we randomly sample 10 sparse ground-truth sets of varying densities with the following sampling strategies:

- **Random Sampling:** Randomly sample a set number of points from the ground-truth depth map.
- **Cluster Based Sampling:** Randomly sample points in clusters from the ground-truth depth map.
- **Pixel Likelihood Sampling:** Each pixel location in the ground-truth depth map is included with a fixed probability.

These sampling strategies allow for the ensemble training to account for a wide range of potential sparse ground-truth sets. Details about the ensemble architecture can be seen in Fig. 2.

To allow for the ensemble’s output to be unconstrained, each model outputs the log of the depth at each pixel location. Each model is then trained to minimize the scale-invariant loss between the predicted depth and the ground-truth depth:

$$\frac{1}{N} \sum_i (\log d_i - \log \hat{d}_i)^2 - \frac{1}{2N^2} \left(\sum_i (\log d_i - \log \hat{d}_i) \right)^2 \quad (7)$$

where N is the total number of depth values and d_i, \hat{d}_i are the ground-truth and predicted depth values respectively.

We train over 50 epochs using an Adam optimizer with a learning rate of 3×10^{-4} . For each ensemble member we use the weights that demonstrate the lowest validation loss.

C. Experimental Setup

We first evaluate our approach on the simulated dataset. For each point in the simulated dataset, the MDE ensemble makes an initial prediction without access to any sparse ground-truth information. We then compare the performance of our method to several baseline point selection strategies.

To validate our framework on real-world data, we perform experiments using ex vivo porcine tracheal models with tumor surrogates using the Virtuoso Endoscopy System as seen in Fig. 3. The system consists of a minimally invasive robot with concentric tube arms deployed through the working channel of an affixed rigid endoscope, pictured in Fig. 1.

Whereas in the simulated setting, a dense ground-truth depth map was readily available, this information does not exist in the physical system. Consequently, depth estimates are obtained through a touch-sensing tool integrated at the distal end of the concentric tube arm. The touch-sensing tool, first introduced in [4], employs a capacitance circuit coupled to an electrosurgery probe running along the arm. Tissue contact produces a sharp change in capacitance, which can be measured in real time, enabling the arm to detect contact events. By using this tool to physically ‘probe’ image-space points, we are able to recover their corresponding depths via robot kinematics.

During each trial, an RGB image of the surgical scene is acquired, and a set of desired measurement locations are selected using one of the point selection strategies (ours or a baseline for comparison). This supplies a list of pixel coordinates to probe for depth information. We use the

calibrated camera intrinsics to back project each of the pixel coordinates as 3D rays whose origin is the camera position. The concentric tube arm, equipped with the touch-sensing tool, will iteratively trace along each of the rays until contact with tissue is sensed. The cartesian coordinate of this ‘sensed point’ can then be used as the known depth information.

Although the ray traced was calculated from the back projection of a specified pixel location, small errors may accrue from several sources. Thus, we re-project the sensed cartesian point back to a pixel coordinate on the image plane. The end result is a list of pixel coordinates with associated depth information. The MDE ensemble then receives the RGB image along with the set of selected depths and predicts a dense depth map of the scene.

D. Baselines

We compare our method against the following baseline point selection strategies:

- **Random Selection:** Points are sampled uniformly at random across the image.
- **Greedy Variance:** The set of pixels with the highest ensemble predictive variance is selected.
- **Stein Variance:** SVGD particles are used over the ensemble variance map to select points.
- **Greedy Gradient:** The set of pixels with the largest magnitude in the gradient of the ensemble’s uncertainty function is selected.
- **Ours (ProbeMDE):** Our proposed cost-aware SVGD method, where SVGD particles evolve over the normalized negated gradient of the ensemble variance.

Each method is evaluated with 5 points being added to the sparse ground-truth set at each iteration. For a fair comparison, each method is run over five iterations (total 25 points in the final prediction). As an additional baseline, we compare the reconstruction quality to a single MDE trained only on RGB image data (no sparse-ground truth).

E. Evaluation Metrics

To evaluate the reconstruction quality of each method, we report standard monocular depth estimation metrics widely used in prior work [34], [35], [36], [37]:

- **Absolute Relative Error (AbsRel):** The mean absolute difference between predicted and ground-truth depth values, normalized by the ground truth: $\frac{1}{N} \sum_i \frac{|d_i - \hat{d}_i|}{d_i}$. This metric is sensitive to proportional errors across depth ranges.
- **Squared Relative Error (SqRel):** The mean squared difference between predictions and ground truth, normalized by the ground truth: $\frac{1}{N} \sum_i \frac{(d_i - \hat{d}_i)^2}{d_i^2}$. This penalizes large outliers more heavily than AbsRel.
- **RMSE:** The square root of the mean squared error between predicted and ground-truth depths: $\sqrt{\frac{1}{N} \sum_i (d_i - \hat{d}_i)^2}$. This metric emphasizes overall prediction accuracy.
- **Log RMSE:** The RMSE computed in log space, which reduces the influence of absolute scale and

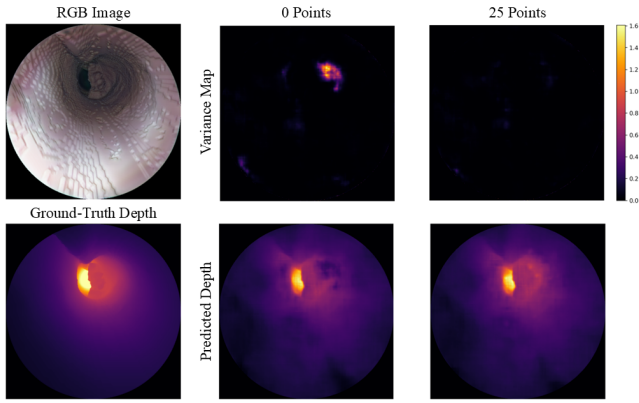


Fig. 4: Qualitative results of ProbeMDE on a simulated central airway obstruction surgical scene. The left column shows the RGB image used as input to the ensemble as well as the ground-truth depth map that corresponds to the image. The middle and right columns show the ensemble’s predictive variance along with the predicted depth map at 0 and 25 propriocepted depth measurements respectively. Qualitatively, the predicted depth maps get refined locally near propriocepted locations to improve detail in the depth map. We also see that as more points are propriocepted, there is a significant reduction in predictive variance across the depth map.

highlights relative differences in depth predictions:

$$\sqrt{\frac{1}{N} \sum_i (\log d_i - \log \hat{d}_i)^2}.$$

- **Scale-Invariant Error (ScInv):** A scale-invariant logarithmic loss that evaluates the consistency of depth relationships regardless of global scale shifts:

$$\frac{1}{N} \sum_i (\log d_i - \log \hat{d}_i)^2 - \frac{1}{2N^2} \left(\sum_i (\log d_i - \log \hat{d}_i) \right)^2$$

- **Threshold Accuracy:** Fraction of pixels such that $\max(\frac{\hat{d}_i}{d_i}, \frac{d_i}{\hat{d}_i}) < 1.25^t$ for $t \in (1, 2, 3)$. These thresholds evaluate prediction accuracy within progressively wider bounds.

Here, d_i and \hat{d}_i denote the ground-truth and predicted depth values, respectively, and N is the total number of depth values in the depth map.

VI. RESULTS

A. Performance on Simulated Data

Table I summarizes the results on our simulated dataset across all baselines and our method. Incorporating a small number of propriocepted ground-truth depths significantly improves depth predictions compared to solely relying on the RGB image. Specifically, the No Sparse Ground-Truth baseline exhibits substantially higher error across all metrics (e.g., SqRel = 2.657 and RMSE = 6.559) compared to any of the active-sensing strategies (e.g., ProbeMDE which achieved SqRel = 0.708 and RMSE = 3.920).

Among the active-sensing baselines, Random Selection already provides large gains, highlighting the value of even unstructured sparse propriocepted measurements. Greedy Variance and Stein Variance also provide significant gains over the no-depth method, but in this setting both are slightly worse than Random Selection. This suggests that targeting




Fig. 5: Qualitative results of ProbeMDE on a physical central airway obstruction surgical scene. The left column shows the RGB image used as input to the ensemble as well as the ground-truth depth map that corresponds to the image. As we do not have access to the full ground-truth depths in the case of our physical experiment, the ground-truth depth map shown is a Voronoi diagram constructed using known depth values at specific points within the image. The middle and right columns show the ensemble’s predictive variance along with the predicted depth map at 0 and 25 propriocepted depth measurements respectively.

regions of high variance magnitude alone is not the most effective heuristic for guiding proprioceptive measurements.

The Greedy Gradient approach yields notable improvements, achieving the lowest RMSE (3.878) and tied for the best accuracy at $\delta < 1.25^3$ (0.981). However, our proposed ProbeMDE approach achieves the strongest overall performance, outperforming all baselines in all other metrics. This demonstrates that utilizing uncertainty gradients and optimizing proprioceptive measurement locations via SVGD provides a more principled and effective approach than variance-only or greedy selection methods.

Fig. 4 shows a qualitative example of ProbeMDE on a simulated environment. In the case where there are no points in the sparse ground-truth set, we see regions of high variance in the variance map. As points are added to the sparse ground-truth set, we see a significant reduction in variance across the prediction while also improving the quality of the prediction, specifically in high depth regions.

Overall, these results confirm that actively guiding proprioception with our cost-aware Stein Gradient method leads to the largest performance improvements in monocular depth estimation accuracy over other active-sensing approaches.

B. Physical Experiments

We further evaluate our framework on physical data collected from surgical phantoms. In the absence of dense ground-truth for the physical case, we exhaustively probe the endoscopic scene to gather a sparse depth map on the order of 300 points which we then use for evaluation. Each active sensing method selects a set of candidate query points which are projected to the nearest valid ground-truth location not already included in the input sparse ground-truth set, ensuring fairness across all selection strategies.

The quantitative results are summarized in Table II. Random Selection achieves the lowest AbsRel (0.242) and high-

Method	AbsRel	SqRel	RMSE	Log RMSE	ScInv	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
No Sparse Ground-Truth	0.352	2.657	6.559	0.376	0.328	0.580	0.832	0.920
Random Selection	0.167	0.739	3.991	0.214	0.188	0.785	0.946	0.978
Greedy Variance	0.171	0.789	4.122	0.221	0.195	0.777	0.944	0.977
Stein Variance	0.170	0.789	4.182	0.221	0.195	0.776	0.942	0.976
Greedy Gradient	0.164	0.720	3.878	0.209	0.183	0.782	0.950	0.981
ProbeMDE (Ours)	0.162	0.708	3.920	0.206	0.180	0.791	0.952	0.981

TABLE I: Performance comparison of active proprioception strategies on simulated data. Results are reported across standard depth estimation metrics. Lower is better for AbsRel, SqRel, RMSE, Log RMSE, and ScInv. Higher is better for $\delta < 1.25^l$. Best results in each column are bolded. We see that ProbeMDE performs consistently the strongest across the majority of metrics, highlighting our approach’s ability to efficiently improve MDE performance.

Method	AbsRel	SqRel	RMSE	Log RMSE	ScInv	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Random Selection	0.242	3.745	14.254	0.391	0.331	0.452	0.809	0.939
Greedy Variance	0.261	3.695	13.797	0.385	0.316	0.377	0.774	0.944
Stein Variance	0.260	3.812	13.588	0.383	0.316	0.406	0.785	0.943
Greedy Gradient	0.264	4.021	14.503	0.410	0.341	0.392	0.767	0.936
ProbeMDE (Ours)	0.250	3.682	13.582	0.380	0.315	0.432	0.791	0.940

TABLE II: Performance comparison of active proprioception strategies on physical data. Results are reported across standard depth estimation metrics. Lower is better for AbsRel, SqRel, RMSE, Log RMSE, and ScInv. Higher is better for $\delta < 1.25^l$. Best results in each column are bolded. These results illustrate how ProbeMDE performs consistently well across the majority of metrics, achieving the best performance in 4 metrics and second best in 3 other metrics.

est $\delta < 1.25$ (0.452), indicating that selecting a set of uniform diverse points can be beneficial in this scenario. However, its overall performance is inconsistent across metrics with a relatively high RMSE (14.254), Log RMSE (0.391), and ScInv (0.331). Greedy and Stein Variance both improve $\delta < 1.25^3$, but both methods perform poorly in most other metrics. The Greedy Gradient method performs the worst overall, showing relatively high errors across all metrics.

Our proposed ProbeMDE method achieves consistently strong performance across metrics, having the best SqRel (3.682), RMSE (13.582), Log RMSE (0.380), and ScInv (0.315) while also producing relatively low errors in the other metrics (second best in AbsRel, $\delta < 1.25$, and $\delta < 1.25^2$). The improvements that ProbeMDE provides across the majority of metrics highlight its robustness in the physical setting where ground-truth supervision is sparse.

Fig. 5 shows an example of ProbeMDE on our physical experimental setup. Similarly to the simulated setting, we see a significant reduction in predictive variance as we increase the number of propriocepted points. Depth predictions also become locally more refined as we propriocept more points in the image.

VII. LIMITATIONS AND FUTURE WORK

While our approach demonstrates clear improvements in monocular depth estimation through uncertainty-guided active sensing, several limitations remain. For instance, our method assumes a static environment, whereas the endoscopic scene will change as the surgery is conducted. In the future, we plan to account for temporal dynamics by considering the historical sequence of depth measurements and endoscopic images. Another limitation of relying heavily on robot probes is the issue of kinematic reachability. Since not every point in the endoscopic image is reachable by the robot, we are unable to obtain depth measurements in

some areas. Future work could explore tracked repositioning of the surgical robot base to enable probing areas which were previously unreachable. In this work we demonstrate that a U-Net trained for a targeted endoscopic procedure undergoes major performance gains with the fusion of sparse proprioceptive measurements. Future work in MDE could benefit from using our proposed framework to incorporate sparse depth measurements into larger foundation models for improved absolute depth accuracy both in surgery and across many domains.

VIII. CONCLUSION

We introduced ProbeMDE, a cost-aware active depth sensing framework that integrates sparse touch-based depth measurements with monocular depth estimation (MDE). By leveraging an ensemble of U-Net depth predictors, we quantify predictive uncertainty via ensemble variance and derive gradients that identify maximally informative locations for new proprioceptive measurements. Using Stein Variational Gradient Descent (SVGD), we select a diverse set of informative probing locations, enabling principled and efficient measurement acquisition.

Through both simulated and physical experiments, we demonstrate that ProbeMDE consistently improves depth estimation accuracy compared to baseline point selection strategies including random sampling, greedy heuristics, and gradient-based heuristics. Notably, our approach maximizes depth accuracy while minimizing the number of used proprioceptive measurements, highlighting the value of uncertainty-guided active sensing in cost-constrained robotic perception.

In future work, we plan to extend ProbeMDE to dynamic surgical environments, incorporating temporal consistency across video frames. Overall, our results illustrate that actively leveraging uncertainty to guide sparse measurements

offers a scalable path toward safe and accurate perception in safety-critical robotics applications.

REFERENCES

- [1] Y. Ming, X. Meng, C. Fan, and H. Yu, "Deep learning for monocular depth estimation: A review," *Neurocomputing*, vol. 438, pp. 14–33, 2021.
- [2] R. Daher, F. Vasconcelos, and D. Stoyanov, "Shades: self-supervised monocular depth estimation through non-lambertian image decomposition," *International Journal of Computer Assisted Radiology and Surgery*, pp. 1–9, 2025.
- [3] J. M. Ferguson, E. B. Pitt, A. A. Ramirez, M. A. Siebold, A. Kuntz, N. L. Kavoussi, E. J. Barth, S. D. Herrell, and R. J. Webster, "Toward practical and accurate touch-based image guidance for robotic partial nephrectomy," *IEEE transactions on medical robotics and bionics*, vol. 2, no. 2, pp. 196–205, 2020.
- [4] N. S. Kumar, J. A. Shrand, S. Sharma, S. D. HerrellIII, N. L. Kavoussi, A. Ghazi, E. J. Barth, I. Oguz, A. Kuntz, and R. J. WebsterIII, "Touching is believing: Contact sensing to enhance situational awareness for concentric tube robots," in *The Hamlyn Symposium on Medical Robotics*, pp. 25–26, 2025.
- [5] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, "Depth Anything V2," *Advances in Neural Information Processing Systems*, vol. 37, pp. 21875–21911, Dec. 2024.
- [6] B. Ke, A. Obukhov, S. Huang, N. Metzger, R. C. Daudt, and K. Schindler, "Repurposing Diffusion-Based Image Generators for Monocular Depth Estimation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9492–9502, 2024.
- [7] L. Shao, L. Bai, C. Du, and X. Chen, "EndoMUST: Monocular Depth Estimation for Robotic Endoscopy via End-to-end Multi-step Self-supervised Training," June 2025. arXiv:2506.16017 [cs].
- [8] M. Tan, Y. Tao, B. Zheng, G. Xie, Z. Xia, and J. Xiong, "Accurate and fast monocular endoscopic depth estimation of structure-content integrated diffusion," *Computerized Medical Imaging and Graphics*, vol. 125, p. 102640, Oct. 2025.
- [9] H. Li, D. Lu, J. d'Almeida, D. Isik, E. K. Aghdam, N. DiSanto, A. Acar, S. Sharma, J. Y. Wu, R. J. Webster III, and I. Oguz, "Monocular absolute depth estimation from endoscopy via domain-invariant feature learning and latent consistency," Nov. 2025. arXiv:2511.02247 [cs].
- [10] B. H. Jeong, H. K. Kim, and Y. D. Son, "Depth estimation from monocular endoscopy using simulation and image transfer approach," *Computers in Biology and Medicine*, vol. 181, p. 109038, Oct. 2024.
- [11] 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 Transactions on Medical Imaging*, vol. 39, pp. 1438–1447, May 2020.
- [12] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, Springer, 2015.
- [13] S. Lee, S. Purushwalkam, M. Cogswell, D. Crandall, and D. Batra, "Why m heads are better than one: Training a diverse ensemble of deep networks," *arXiv preprint arXiv:1511.06314*, 2015.
- [14] D. Milanés-Hermosilla, R. Trujillo Codorniu, R. López-Baracaldo, R. Sagaró-Zamora, D. Delisle-Rodríguez, J. J. Villarejo-Mayor, and J. R. Núñez-Álvarez, "Monte carlo dropout for uncertainty estimation and motor imagery classification," *Sensors*, vol. 21, no. 21, p. 7241, 2021.
- [15] R. Camarasa, D. Bos, J. Hendrikse, P. Nederkoorn, E. Kooi, A. Van Der Lugt, and M. De Bruijne, "Quantitative comparison of monte-carlo dropout uncertainty measures for multi-class segmentation," in *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis: Second International Workshop, UNSURE 2020, and Third International Workshop, GRAIL 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 8, 2020, Proceedings 2*, pp. 32–41, Springer, 2020.
- [16] Y. Mae, W. Kumagai, and T. Kanamori, "Uncertainty propagation for dropout-based bayesian neural networks," *Neural Networks*, vol. 144, pp. 394–406, 2021.
- [17] G. Wang, W. Li, M. Aertsen, J. Deprest, S. Ourselin, and T. Vercauteren, "Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks," *Neurocomputing*, vol. 338, pp. 34–45, Apr. 2019.
- [18] Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. Dillon, B. Lakshminarayanan, and J. Snoek, "Can you trust your model' s uncertainty? Evaluating predictive uncertainty under dataset shift," in *Advances in Neural Information Processing Systems*, vol. 32, Curran Associates, Inc., 2019.
- [19] S. Fort, H. Hu, and B. Lakshminarayanan, "Deep ensembles: A loss landscape perspective," *arXiv preprint arXiv:1912.02757*, 2019.
- [20] T. Pearce, A. Brintrup, M. Zaki, and A. Neely, "High-quality prediction intervals for deep learning: A distribution-free, ensemble approach," in *International conference on machine learning*, pp. 4075–4084, PMLR, 2018.
- [21] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, vol. 30, 2017.
- [22] M. Xiang, J. Zhang, N. Barnes, and Y. Dai, "Measuring and Modeling Uncertainty Degree for Monocular Depth Estimation," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 34, pp. 5716–5727, July 2024.
- [23] R. Marsal, F. Chabot, A. Loesch, W. Grolleau, and H. Sahbi, "Mono-Prob: Self-Supervised Monocular Depth Estimation with Interpretable Uncertainty," in *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 3625–3634, Jan. 2024. ISSN: 2642-9381.
- [24] S. Landgraf, R. Qin, and M. Ulrich, "A Critical Synthesis of Uncertainty Quantification and Foundation Models in Monocular Depth Estimation," Jan. 2025. arXiv:2501.08188 [cs].
- [25] N. Chong and F. Yang, "A monocular medical endoscopic images depth estimation method based on a confidence-guided dual-branch siamese network," *Biomedical Signal Processing and Control*, vol. 102, p. 107123, Apr. 2025.
- [26] T. Assaf, C. Roke, J. Rossiter, T. Pipe, and C. Melhuish, "Seeing by touch: Evaluation of a soft biologically-inspired artificial fingertip in real-time active touch," *Sensors*, vol. 14, no. 2, pp. 2561–2577, 2014.
- [27] Y. B. Kim, G. Kang, G. K. Yee, A. Kim, W. S. You, Y. H. Lee, F. Liu, H. Moon, J. C. Koo, and H. R. Choi, "Exploration and reconstruction of unknown object by active touch of robot hand," *Intelligent Service Robotics*, vol. 8, no. 3, pp. 141–149, 2015.
- [28] U. Martinez-Hernandez, T. J. Dodd, M. H. Evans, T. J. Prescott, and N. F. Lepora, "Active sensorimotor control for tactile exploration," *Robotics and Autonomous Systems*, vol. 87, pp. 15–27, 2017.
- [29] A.-H. Shahidzadeh, S. J. Yoo, P. Mantripragada, C. D. Singh, C. Fermüller, and Y. Aloimonos, "Actexplor: Active tactile exploration on unknown objects," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3411–3418, IEEE, 2024.
- [30] J. Xu, S. Song, and M. Ciocarlie, "Tandem: Learning joint exploration and decision making with tactile sensors," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10391–10398, 2022.
- [31] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose bayesian inference algorithm," *Advances in neural information processing systems*, vol. 29, 2016.
- [32] T. Park, A. A. Efros, R. Zhang, and J.-Y. Zhu, "Contrastive learning for unpaired image-to-image translation," in *European conference on computer vision*, pp. 319–345, Springer, 2020.
- [33] Virtuoso Surgical, Nashville, TN, *Virtuoso Endoscopy System*. <https://virtuososurgical.net/technology/>. Accessed 2025-09-15.
- [34] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Advances in neural information processing systems*, vol. 27, 2014.
- [35] D. Eigen and R. Fergus, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," in *Proceedings of the IEEE international conference on computer vision*, pp. 2650–2658, 2015.
- [36] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 270–279, 2017.
- [37] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab, "Deeper depth prediction with fully convolutional residual networks," in *2016 Fourth international conference on 3D vision (3DV)*, pp. 239–248, IEEE, 2016.