

UC San Diego

UC San Diego Previously Published Works

Title

KineDepth: Utilizing Robot Kinematics for Online Metric Depth Estimation

Permalink

<https://escholarship.org/uc/item/3jm3n4t5>

Journal

2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 00

Authors

Atar, Soofiyan

Zhi, Yuheng

Richter, Florian

et al.

Publication Date

2025-10-25

DOI

10.1109/iros60139.2025.11247739

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nd/4.0/>

Peer reviewed

KineDepth: Utilizing Robot Kinematics for Online Metric Depth Estimation

Soofiyan Atar, Yuheng Zhi, Florian Richter, and Michael Yip

Abstract—Depth perception is essential for a robot’s spatial and geometric understanding of its environment, with many tasks traditionally relying on hardware-based depth sensors like RGB-D or stereo cameras. However, these sensors face practical limitations, including issues with transparent and reflective objects, high costs, calibration complexity, spatial and energy constraints, and increased failure rates in compound systems. While monocular depth estimation methods offer a cost-effective and simpler alternative, their adoption in robotics is limited due to their output of relative rather than metric depth, which is crucial for robotics applications. In this paper, we propose a method that utilizes a single calibrated camera, enabling the robot to act as a “measuring stick” to convert relative depth estimates into metric depth in real-time as tasks are performed. Our approach employs an LSTM-based metric depth regressor, trained online and refined through probabilistic filtering, to accurately restore the metric depth across the monocular depth map, particularly in areas proximal to the robot’s motion. Experiments with real robots demonstrate that our method significantly outperforms current state-of-the-art monocular metric depth estimation techniques, achieving a 22.1% reduction in depth error and a 52% increase in success rate for a downstream task.

I. INTRODUCTION

Robot manipulators are quickly becoming a commodity technology that can carry out laborious human tasks, such as household cleaning, food preparation and dishwashing, pantry stocking items, etc [1]–[3]. The applications are broad, from home-care robot assistants [4] to robots that operate and assist in the operating room [5]. In interacting with the world, the positional accuracy of the robot and its estimate of targets in the environment are basic necessities such that robot end-effectors can accurately and precisely reach, grasp, manipulate, and potentially deform objects and/or environments.

Traditionally, 3D sensors such as RGB-D cameras and stereo-cameras were used to acquire geometrical and spatial information in the workspace and of objects of interest [6]. As a standard approach, robotic manipulation systems with depth sensors have been proven to work well in a wide variety of scenarios [7]–[11].

However, there are benefits and limitations to 3D sensors. Technologies such as structured light, stereo vision, or LiDAR, can have readings thrown off by reflective surfaces, transparent objects, ambient light interference, and limited textural features [12]. They also tend to have more defined




Fig. 1: Visualization of point clouds for different depth estimation methods applied to a target object (cup). The orange point cloud represents KineDepth, the gray shows scaled relative depth from DepthAnything v2 [21], and the purple represents ground truth metric depth. In the zoomed-in view, KineDepth closely aligns with the ground truth, while the relative depth shows significant local and global misalignment.

and limited operating ranges that have diminishing accuracy at both close and far distances [13].

Monocular depth estimation (MDE) methods have more recently become available, where depth may be estimated from a single RGB image and a learned model [14]. This alleviates some issues associated with depth sensors and makes monocular RGB with a learned depth model an interesting and more widely accessible alternative. In support of significant potential in MDE, very recently, training of foundation models with large-scale datasets has shown significant leaps in resolving depth estimates across a highly diverse range of environments with a single model [15], [16].

Despite the advances in MDEs, a main challenge of their use in robotic control applications is that MDEs typically yield relative depth estimates, which, while useful, suffer from some practical drawbacks, such as having no sense of the scale of a scene. While *metric* monocular depth estimators (MMDEs) are very recently becoming available and provide absolute depth estimates [17]–[20], their estimates are not as stable when used live, can vary frame to frame under dynamic environments, and are still susceptible to inaccuracies due to noise in camera parameters. This inaccuracy creates real problems for robots with solving for the correct inverse kinematics, image Jacobians, and evaluating workspace reachability for driving a robot end effector to a target. Thus, this limits the ability of the

All the authors are with the Electrical and Computer Engineering department, University of California San Diego.

This work was supported by the US Army Telemedicine and Advanced Technologies Research Center and NSF Career Award #2045803.

robot to do manipulation tasks requiring higher accuracy and precision.

To this end, we propose a *KineDepth*, a novel methodology that derives accurate metric depth from monocular RGB images by incorporating the reference geometry of a robotic manipulator. Consider your robot as a constant measuring stick in the scene, with its known geometry – it can be used as a local reference for metric depth estimation. Central to our approach is the integration of two aspects: (1) an off-the-shelf depth estimator such as Depth Anything V2 [21] that provides relative depth estimation using monocular cameras, without requiring fine-tuning, and (2) a novel *depth scale regressor* — a mathematical function that maps relative depth to metric depth that dynamically adapts based on the manipulator’s kinematic chain as it approaches a target in the task space. As the manipulator moves closer to the target, the precision of the reference geometry increases, thereby enhancing the accuracy of metric depth estimation.

With the ultimate goal of working in novel, unstructured environments, we focus the KineDepth approach on working online and without pre-training. The key contributions of the paper are summarized as follows:

- **Novel Monocular Metric Depth Estimation Framework:** We present a new MMDE framework that synthesizes relative depth from a foundation model with observed robotic manipulator geometry.
- **Dynamic Depth Scale Regressor:** We introduce a dynamic depth scale regressor that adapts to the manipulator’s kinematic chain, enhancing depth precision as the end-effector approaches the target.
- **Online Tracking and Filtering:** We provide KineDepth with online estimation and filtering so that it may be deployed without pretraining and, therefore, for unstructured and unseen environments. It leverages an LSTM network for temporal modeling, which predicts the depth scale regressor refined by a Kalman filter for consistency and noise reduction.

Moreover, through extensive evaluations of random trajectories and pick-and-place tasks, we demonstrate that our approach outperforms state-of-the-art models in depth prediction accuracy, providing a scalable solution using only monocular cameras for high-precision manipulation tasks.

II. RELATED WORK

Monocular depth estimation has been a pivotal area of research in computer vision, particularly for its applications in robotics and 3D reconstruction. Traditional methods for Single-Image Depth Estimation (SIDE) [17]–[20], [22], [23] can be categorized into metric depth regression and relative depth estimation [15], [16], [21], [24], [25]. Metric depth models often suffer from overfitting due to training on singular datasets, leading to poor generalization in unseen environments or varying depth ranges. Relative depth models, trained using scale-invariant losses on diverse datasets, tend to generalize better but lack the metric scale necessary for applications requiring absolute depth measurements. Recent efforts aim to recover metric information by combining

monocular depth estimation with additional modules or by reformulating depth regression as a classification task.

Advancements in distribution learning have led to methods that treat depth estimation as a combined classification-regression problem, reasoning about depth value distributions across images. Approaches like AdaBins [26], LocalBins [27], and PixelFormer [28] have improved depth estimation by adaptively discretizing predicted depth ranges into bins, allowing for more precise depth predictions.

Monocular Metric Depth Estimation (MMDE) using neural networks was pioneered by transformer-based architectures [29]. Despite significant progress on benchmarks like NYU Depth v2 [30] and KITTI [31], MMDE models often struggle with zero-shot generalization due to domain shifts in appearance and geometry.

To enhance generalization across diverse domains, recent methods incorporate camera awareness by integrating external camera parameters or normalizing outputs based on intrinsic properties [18], [32], [33]. However, these approaches typically rely on noiseless camera intrinsics and predefined back-projection operations, limiting their applicability, especially in the absence of accurate camera information.

The use of large-scale data has propelled advancements in depth estimation. Models trained on extensive datasets [34]–[36] demonstrate improved generalization. Self-supervised learning approaches [37], [38] and large-scale pretraining [39], [40] have further enhanced model robustness. These methods are prone to domain shift in new environments, and they may overfit to dataset-specific biases while lacking the fine-grained accuracy needed for precise depth estimation, issues we aim to address in this paper.

Despite the progress, there remains a gap in achieving accurate metric depth estimation from monocular RGB images without relying on additional sensors or prior knowledge. Existing methods either generalize poorly to new environments or lack the metric scale necessary for precise applications in robotic manipulation and 3D reconstruction.

III. METHODS

This section outlines the pipeline for Monocular Metric Depth Estimation (MMDE) as shown in Fig. 2. At the core of this method is updating a depth regressor in an online manner, which maps relative depth to metric depth using the visual measurements of the robot manipulator with known kinematics as known geometry.

A. Depth Regressor

The depth regressor, $Z(u, v) = f(r(u, v), \beta)$ for pixel location $[u, v]^T \in \mathbb{R}^2$, in our system is a function, $f(\cdot)$, that converts the relative depth image, $r(\cdot)$, obtained from a monocular depth estimation model, such as Depth Anything v2 model [21], into metric depth, $Z(\cdot)$. The depth regressor is parameterized by β which we estimate in an online fashion using the robot as a geometric reference. We constrained the depth regressor, $f(\cdot)$, to simple models to ensure stability when applying the regressor to downstream control tasks and evaluated various functions, including power law, rational,




Fig. 2: Our framework leverages the robot as a geometric reference to convert relative depth from monocular depth estimation techniques such as Depth Anything V2 to metric depth. The geometric references are keypoints on the robot, γ , which are applied to update our LSTM and Kalman Filters to estimate our depth regressor parameters, β . The depth regressor parameters create a polynomial fit to finally convert the relative depth to metric depth. The metric depth is used as feedback in downstream robotic control and manipulation tasks such as end-effector control.

Gaussian, logarithmic, linear, and polynomial regressors. We ran an experiment on fitting Depth Anythingv2 predictions to ground truth using data collected from our experiments and SciPy's *optimize.fit* function. The polynomial regressor demonstrated the highest accuracy in fitting, as shown in Tab. I.

Thus, our depth regressor is defined as

$$Z(u, v) = \beta_2 r(u, v)^2 + \beta_1 r(u, v) + \beta_0. \quad (1)$$

where coefficients $\beta = [\beta_2, \beta_1, \beta_0]^\top$ are polynomial coefficients that will be estimated in an online fashion.

B. Online Estimation

We estimate the depth regressor parameters, β , in an online fashion using the robot as a geometric reference. The robot geometric references are defined as keypoints, $\gamma^i \in \mathbb{R}^2$, which correspond to the following observed metric depth, o^i , from the robot kinematics

$$o^i = \left[\mathbf{T}_c^b \prod_{i=1}^n \mathbf{T}_{i-1}^i(\theta_i) p^i \right]_z \quad (2)$$

where $\mathbf{T}_c^b \in SE(3)$ is the camera to base transform, $\mathbf{T}_{i-1}^i(\theta_i) \in SE(3)$ is the i -th joint transform with joint angle θ_i , and p^i is the translational offset of the 3D keypoint.

TABLE I: Table shows the accuracy fit for different regressor to fit the relative depth to estimate the metric depth

Regressor	Fit Percentage
Gaussian Model	26.77%
Logarithmic Model	81.61%
Power-law Model	85.05%
Rational Model	85.09%
Linear model	88.42%
Polynomial (2nd order) model	95.30%

We assume knowledge of the extrinsics, $\mathbf{T}_c^b$ [41], robot kinematics $\mathbf{T}_{i-1}^i(\theta_i)$, and keypoint offsets p^i to compute geometric references, o^i , for our online estimation of the depth regressor parameters, β . Note o^i should correspond to the metric depth at the keypoint location, $Z(\gamma^i)$. The keypoints, γ_t^i , will be detected using CoTracker [42] at every timestep t , however other keypoint estimation techniques could be used [43]. We present three different approaches for the online estimation: Kalman Filter, LSTM, and a Hybrid approach. All three approaches are outlined in Alg. 1.

1) *Kalman Filter*: To utilize a Kalman Filter for online parameter estimation of the depth regressor parameters, β , we define initialization, motion, and observation models. The parameters are initialized as a Gaussian with a linear scale, and the motion model is assumed to be mean-zero and Gaussian since we do not expect the parameters to change drastically between timesteps. Mathematically, the initialization and motion models are defined as

$$\beta_0 \sim \mathcal{N}([0, 1, 0]^\top, \Sigma_0) \quad \beta_t \sim \mathcal{N}(\beta_{t-1}, \Sigma_t) \quad (3)$$

where Σ_0 is the initial covariance and Σ_t is the motion model covariance.

Finally, the observation model is based on our reference geometry observations, o^i defined in Eq. 2, and modelled as

$$o_t^i \sim \mathcal{N}(\mathbf{H}\beta_t, \Sigma_o) \quad (4)$$

where $\mathbf{H} = [r(\gamma_t^i)^2, r(\gamma_t^i), 1]$ and Σ_o is the observation model covariance. With all the models linear and Gaussian, the Kalman Filter can be employed to estimate β .

2) *LSTM*: Our secondary estimation approach uses an LSTM which is trained in an online fashion to estimate the depth regressor parameters, β . The LSTM takes as input the relative depth at each keypoint and predicts hidden state, $\mathbf{h}_t$:

$$\mathbf{h}_t, \mathbf{c}_t = \text{LSTM}([r(\gamma_t^1), \dots, r(\gamma_t^M)]^\top, \mathbf{h}_{t-1}, \mathbf{c}_{t-1}) \quad (5)$$

where $\mathbf{c}_t$ is the copy state. The LSTM is followed by a multilayer perceptron (MLP) to ultimately predict the depth regressor coefficients, β_t , and denoted as

$$\beta_t = \text{MLP}(\mathbf{h}_t) \quad (6)$$

Similar to the Kalman Filter technique, we train the model using the depth observations, o^i , with the following loss

$$\sum_{i=1}^M \mathcal{L}_h(f(r(\gamma_t^i), \beta_t), o_t^i) + \mathcal{L}_{aux} \quad (7)$$

where $\mathcal{L}_h$ is the Huber loss with threshold δ ,

$$\mathcal{L}_h(y, \hat{y}) = \begin{cases} \frac{1}{2}(y - \hat{y})^2 & \text{for } |y - \hat{y}| \leq \delta \\ \delta (|y - \hat{y}| - \frac{1}{2}\delta) & \text{otherwise} \end{cases} \quad (8)$$

$f(\cdot)$ is the depth regressor evaluated at the keypoint location of the relative depth map, $r(\gamma_t^i)$, with parameters β_t predicted by the LSTM and MLP and $\mathcal{L}_{aux}$ is an auxiliary loss to help with training. The auxiliary loss is applied to auxiliary outputs of the model: unique polynomial coefficients for each keypoint, β_t^i , and the metric depth at every keypoint, m_t^i . Since the metric depth should be inferred from the polynomial coefficients, an additional MLP is applied to the head to predict m_t^i . The auxiliary loss is defined as

$$\mathcal{L}_{aux} = \alpha_1 \sum_i^M \mathcal{L}_h(f(r(\gamma_t^i), \beta_t^i), o_t^i) + \alpha_2 \sum_i^M \mathcal{L}_h(m_t^i, o_t^i) \quad (9)$$

where α_1, α_2 are weightings. Furthermore, we append the last N relative depth values at each keypoint to the input of the LSTM, i.e. $r(\gamma_{t-N}^i), \dots, r(\gamma_t^i)$, and run the LSTM τ steps for one time-step t .

3) *Hybrid*: Experimentally, we found the LSTM estimation approach provided high accuracy but low precision, i.e. large fluctuations in the depth regressor parameters. These fluctuations would make downstream tasks such as robotic manipulation become unstable. To address this issue and improve the consistency of the results, we incorporate a Kalman Filter to smooth the LSTM-generated depth regressor parameters, β_t .

First, we maintain the online training described in the previous LSTM subsection. Secondly, we apply a Kalman Filter by considering the LSTM a learned motion model. The initialization and motion models are then defined as

$$\beta_t \sim \mathcal{N}(\text{MLP}(\mathbf{h}_t), \Sigma_t) \quad (10)$$

where $\mathbf{h}_t$ is the hidden state from Eq. 5, the MLP infers β_t from the hidden state as shown in Eq. 6, and Σ_t is the motion model covariance. Finally, we directly apply the Kalman Filter update step by using the observation model described in Eq. 4.

C. End Effector Control

We employ the metric depth computed from Eq. 1 using the estimated regression parameters, β , to control the end-effectors position in a closed-loop fashion. We assume a goal in pixel space is given, $\mathbf{g}^p$, such as a cup detected in the

Algorithm 1 Online Estimation of β

```

1: if Kalman Filter then
2:    $\beta_0 \leftarrow [0, 1, 0], \mathbf{P}_0 \leftarrow \Sigma_0$   $\triangleright$  Kalman Estimation
3: end if
4: if LSTM or Hybrid then
5:   LSTM, MLP  $\leftarrow$  initializeModel()
6:    $\mathbf{h}_0 \leftarrow 0, \mathbf{c}_0 \leftarrow 0$ 
7: end if
8: for  $t = 1$  to  $T$  do
9:    $I_t \leftarrow \text{getRGBImage}()$ 
10:   $\theta_t \leftarrow \text{getJointAngles}()$ 
11:   $r_t \leftarrow \text{monocularDepth}(I_t)$   $\triangleright$  DepthAnythingV2
12:   $\gamma_t \leftarrow \text{keypoint}(I_t)$   $\triangleright$  CoTracker
13:   $o_t \leftarrow \text{observedMetricDepth}(\theta_t)$ 
14:  if Kalman Filter then
15:     $\beta_t, \mathbf{P}_t \leftarrow \text{predictStep}(\beta_{t-1}, \mathbf{P}_{t-1}, \Sigma_t)$ 
16:  else  $\triangleright$  LSTM or Hybrid
17:    for  $k = 1$  to  $\tau$  do
18:       $\mathbf{h}_t \leftarrow \text{LSTM}(r_t(\gamma_t), \mathbf{h}_{t-1}, \mathbf{c}_{t-1})$ 
19:       $\beta_t \leftarrow \text{MLP}(\mathbf{h}_t)$ 
20:       $\mathcal{L} \leftarrow \text{computeLoss}(\beta_t, o_t)$ 
21:      LSTM, MLP  $\leftarrow \text{backPropagate}(\mathcal{L})$ 
22:    end for
23:  end if
24:  if Kalman Filter or Hybrid then
25:     $\beta_t, \mathbf{P}_t \leftarrow \text{updateStep}(\beta_t, \mathbf{P}_t, o_t, \Sigma_o)$ 
26:  end if
27: end for

```

image. The goal is converted to 3D by inverting the pin-hole projection

$$\mathbf{g}_t = f(r(\mathbf{g}^p), \beta_t) \mathbf{K}^{-1} \mathbf{g}^p \quad (11)$$

where $\mathbf{K}$ is the camera intrinsics matrix. While the end-effector could be computed using the robot kinematics, similar to the Eq. 2, we compute the 3D position of the end-effector using the depth regressor in the same manner as Eq. 11 to minimize any potential bias from the estimated goal and end-effector which would lead the robot to miss the goal. Therefore, the state is computed as

$$\mathbf{x}_t = f(r(\gamma_t^x), \beta_t) \mathbf{K}^{-1} \gamma_t^x \quad (12)$$

where γ_t^x is the end-effector keypoint. Finally, we define the action space as a change in end-effector, i.e., $\mathbf{x}_t = \mathbf{x}_{t-1} + \mathbf{u}_t$. A Linear Quadratic Regulator (LQR) is solved at every timestep with the following cost

$$J = \sum_{i=t}^{\infty} \left((\mathbf{x}_i - \mathbf{g}_t)^\top \mathbf{Q} (\mathbf{x}_i - \mathbf{g}_t) + \mathbf{u}_i^\top \mathbf{R} \mathbf{u}_i \right) \quad (13)$$

where $\mathbf{Q}$ and $\mathbf{R}$ are weighting matrices. Only one action is taken after solving the LQR, and then the entire process is repeated (compute goal, end-effector, and LQR) after estimating the latest depth regressor parameters, β_t .

IV. EXPERIMENTS AND RESULTS

A. Implementation details

Our experiment focuses on running KineDepth on Franka Robot, and all the experiments are compared using the same

TABLE II: Comparison of metric depth estimation accuracy across keypoints ($\gamma^1 - \gamma^5$) and overall scene error for different methods within the task space, including both the robotic manipulator and the task environment. The table reports errors for each method, with KineDepth variants (LSTM, KF, and hybrid) and traditional methods such as Metric 3D v2, Zoedepth, and DepthAnything v2.

Method	γ^1 Error	γ^2 Error	γ^3 Error	γ^4 Error	γ^5 Error	Overall scene Error
DepthAnything v2 [21]	1.544 ± 0.014	1.509 ± 0.035	1.608 ± 0.007	1.329 ± 0.148	1.326 ± 0.172	1.423 ± 0.059
UniDepth [18]	1.234 ± 0.017	1.255 ± 0.013	1.188 ± 0.010	1.214 ± 0.043	1.196 ± 0.044	1.242 ± 0.031
Zoedepth [19]	0.896 ± 0.004	0.706 ± 0.008	0.825 ± 0.008	0.885 ± 0.030	0.951 ± 0.028	0.878 ± 0.025
NeWCRF [44]	0.337 ± 0.002	0.515 ± 0.025	0.153 ± 0.004	0.450 ± 0.086	0.361 ± 0.078	0.195 ± 0.043
Metric 3D v2 [17]	0.150 ± 0.004	0.079 ± 0.003	0.044 ± 0.001	0.146 ± 0.046	0.074 ± 0.009	0.068 ± 0.003
KineDepth (LSTM)	0.121 ± 0.004	0.089 ± 0.002	0.059 ± 0.001	0.048 ± 0.001	0.035 ± 0.003	0.099 ± 0.005
KineDepth (KF)	0.103 ± 0.000	0.073 ± 0.002	0.044 ± 0.000	0.060 ± 0.006	0.029 ± 0.001	0.061 ± 0.001
KineDepth (hybrid)	0.035 ± 0.002	0.032 ± 0.001	0.036 ± 0.000	0.050 ± 0.002	0.011 ± 0.000	0.053 ± 0.001

TABLE III: Comparison of metric depth estimation performance for different methods on a marker pick-and-place task involving various scene configurations: easy, intermediate, and hard. The table presents accuracy percentages for each configuration, as well as overall accuracy for each method.

Method	Easy configuration	Intermediate configuration	Hard Configuration	Overall accuracy
ZoeDepth [19]	0.0%	0.0%	0.0%	0.0%
UniDepth [18]	0.0%	0.0%	0.0%	0.0%
DepthAnything v2 [21]	0.0%	0.0%	0.0%	0.0%
NeWCRF [44]	0.0%	0.0%	0.0%	0.0%
Metric 3D v2 [17]	33.3%	25.0%	25.0%	28.0%
KineDepth (LSTM)	33.3%	37.5%	37.5%	36.0%
KineDepth (KF)	81.8%	42.8%	66.7%	64.0%
KineDepth (hybrid)	100%	75%	62.5%	80.0%




Fig. 3: Pick-and-place task: Following a coarse trajectory using monocular depth estimation, where metric depth is calculated via KineDepth and tracked throughout the process.

robot. Fig. 3 depicts the setup of our environment where the camera used is Azure Kinect and the GPU used is Nvidia GTX 1080, which runs the DepthAnything V2 model [21] with the Depth-Anything-V2-Large, CoTracker [42] and online training of LSTM network. The Azure Kinect is used to get the ground truth using the camera, and for the experiments, only the RGB channel was utilized. For the marker pick and place task, we used the setup depicted in Fig. 3, where the marker and cup were randomly placed.

In the LSTM implementation, we configured the key

hyperparameters as follows. The learning rate was set to 1×10^{-2} , with a batch size equal to the number of visible keypoints M in the current frame. The hidden state size was fixed at 128, and the network was trained for τ 20 time steps per update. Σ_0 is a identity matrix, Σ_t are a diagonal matrix with 0.5 and Σ_o are a diagonal matrix with 0.03. In Huber loss, the threshold is set at $\delta = 1.0$, while α_1 and α_2 are set to 1.0. Meanwhile, N is set to 10 to store the trajectory history for the input of the LSTM.

B. Experimental Setup

The setup chosen here is to showcase the accuracy required to grasp the marker and place it in the cup where the goal given to the model was the pixel coordinates and intrinsic parameters. In our experiments, we compared our method against 3 variants of KineDepth, metric 3D v2 [17], Unidepth [18], NewCRFs [44], ZoeDepth [19] and DepthAnything v2 [21] whose decoder is fine-tuned using ZoeDepth architecture. All the experiments are performed on real-world evaluation using the setup explained above.

C. Trajectory Following for MMDE

To evaluate the performance and robustness of metric depth estimation methods, we ran 25 random trajectories for each method and calculated the error at each keypoint, using the depth image as ground truth at the pixel coordinates γ^i . As shown in Tab. II, KineDepth achieved the lowest mean error across all keypoints and the overall scene, which includes the robotic manipulator and task space region. KineDepth (hybrid) had a 5.3 cm overall error, 1.5 cm lower than Metric 3D v2, which outperformed other MMDE methods.




Fig. 4: Physical Experiments: The upper row displays successful pick-and-place executions, while the bottom row illustrates failure cases. The figure covers a range of scene configurations and varying levels of grasping difficulty. Failures are primarily due to incorrect depth estimation or target poses prediction errors, leading to either missed grasps or inaccurate positioning of the marker.




Fig. 5: (a) Trajectory comparison of the end effector across different methods, where the green curve represents the ground truth and the blue curve (KineDepth) shows the closest match. (b) Comparison of the relative depth trajectory with the ground truth, illustrating no correlation between the two.

KineDepth (KF) estimates depth by modeling the depth regressor as the state in a Kalman filter, while KineDepth (LSTM) uses the predicted depth regressor directly from the network. Other methods had errors exceeding 19.5 cm for the overall scene.

D. Marker pick place task

To quantitatively assess the methods, we designed a pick-and-place task with a marker and cup, requiring high precision. As shown in Fig. 3, the setup involves random marker and cup placements, with challenging configurations outlined in Tab. III. Misperceptions often occur when the marker is placed near the camera or diagonally, complicating trajectory planning. Fig. 4 shows different configurations, while Tab. III demonstrates KineDepth’s superiority, with the hybrid version achieving 80% accuracy on 25 pick-place tasks. On harder configurations, KineDepth (KF) slightly outperformed the hybrid version. Other baselines, apart from Metric 3D v2, had no success due to poor metric depth estimates, as supported by results compared in Tab. II.

V. DISCUSSION AND CONCLUSION

As shown in Tab. II and Tab. III, KineDepth outperforms all baselines, with the hybrid version achieving even better

results. Fig. 5(a) illustrates the end effector trajectory for various methods, where KineDepth closely matches the ground truth (depth camera). Other KineDepth variants also exhibit low error, while Metric 3D v2 produces a trajectory that significantly deviates from the ground truth. Furthermore, Fig. 5(b) clearly shows that the relative depth has no correlation with the true depth, indicating that simply scaling the relative depth does not achieve accurate metric depth estimation.

In conclusion, the proposed KineDepth framework introduces a depth regressor that converts relative depth r to the metric depth with high accuracy in the task space region. Further, using Kalman filtering to smooth out the depth regressor predicted from the LSTM network boosted our performance in getting a more accurate estimation. KineDepth delivers improved metric depth accuracy in the task space region, and it achieved 1.5cm less error for the overall task space region as compared to Metric 3D v2. KineDepth (hybrid) provides 22.1% higher accuracy than the state-of-the-art MMDE, achieving the lowest overall scene error of 0.053 m across all key points.

REFERENCES

- [1] M. Suomalainen, Y. Karayiannidis, and V. Kyrki, "A survey of robot manipulation in contact," *Robotics and Autonomous Systems*, vol. 156, p. 104224, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889022001312>
- [2] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annual review of control, robotics, and autonomous systems*, vol. 3, no. 1, pp. 297–330, 2020.
- [3] Y. Jiang, Z. Huang, B. Yang, and W. Yang, "A review of robotic assembly strategies for the full operation procedure: planning, execution and evaluation," *Robotics and Computer-Integrated Manufacturing*, vol. 78, p. 102366, 2022.
- [4] B. Sawik, S. Tobis, E. Baum, A. Suwalska, S. Kropińska, K. Stachnik, E. Perez-Bernabeu, M. Cildoz, A. Agustín, and K. Wiczorowska-Tobis, "Robots for elderly care: Review, multi-criteria optimization model and qualitative case study," *Healthcare*, vol. 11, p. 1286, 04 2023.
- [5] X. Sun, J. Okamoto, K. Masamune, and Y. Muragaki, "Robotic technology in operating rooms: a review," *Current Robotics Reports*, vol. 2, 09 2021.
- [6] W. Wang, J. Li, Y. Zhu, Z. Xu, Z. Che, Y. Peng, C. Shen, D. Liu, F. Feng, and J. Tang, "Visual robotic manipulation with depth-aware pretraining," 2024. [Online]. Available: <https://arxiv.org/abs/2401.09038>
- [7] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," 2022. [Online]. Available: <https://arxiv.org/abs/2209.05451>
- [8] C. Bao, H. Xu, Y. Qin, and X. Wang, "Dexart: Benchmarking generalizable dexterous manipulation with articulated objects," 2023. [Online]. Available: <https://arxiv.org/abs/2305.05706>
- [9] Y. Xu, W. Wan, J. Zhang, H. Liu, Z. Shan, H. Shen, R. Wang, H. Geng, Y. Weng, J. Chen, T. Liu, L. Yi, and H. Wang, "Unidexgrasp: Universal robotic dexterous grasping via learning diverse proposal generation and goal-conditioned policy," 2023. [Online]. Available: <https://arxiv.org/abs/2303.00938>
- [10] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.
- [11] J. Li and D. J. Cappelleri, "Sim-suction: Learning a suction grasp policy for cluttered environments using a synthetic benchmark," *IEEE Transactions on Robotics*, vol. 40, p. 316–331, 2024. [Online]. Available: <http://dx.doi.org/10.1109/TRO.2023.3331679>
- [12] O. Wasenmüller and D. Stricker, "Comparison of kinect v1 and v2 depth images in terms of accuracy and precision," in *Computer Vision – ACCV 2016 Workshops*, C.-S. Chen, J. Lu, and K.-K. Ma, Eds. Cham: Springer International Publishing, 2017, pp. 34–45.
- [13] K. Khoshelham and S. O. Elberink, "Accuracy and resolution of kinect depth data for indoor mapping applications," *Sensors*, vol. 12, no. 2, pp. 1437–1454, 2012. [Online]. Available: <https://www.mdpi.com/1424-8220/12/2/1437>
- [14] Y. Ming, X. Meng, C. Fan, and H. Yu, "Deep learning for monocular depth estimation: A review," *Neurocomputing*, vol. 438, pp. 14–33, 2021.
- [15] S. Patni, A. Agarwal, and C. Arora, "Ecodepth: Effective conditioning of diffusion models for monocular depth estimation," 2024. [Online]. Available: <https://arxiv.org/abs/2403.18807>
- [16] Y. Wang, Y. Liang, H. Xu, S. Jiao, and H. Yu, "Sqldepth: Generalizable self-supervised fine-structured monocular depth estimation," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 6, pp. 5713–5721, Mar. 2024. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/28383>
- [17] M. Hu, W. Yin, C. Zhang, Z. Cai, X. Long, H. Chen, K. Wang, G. Yu, C. Shen, and S. Shen, "Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation," 2024. [Online]. Available: <https://arxiv.org/abs/2404.15506>
- [18] L. Piccinelli, Y.-H. Yang, C. Sakaridis, M. Segu, S. Li, L. V. Gool, and F. Yu, "Unidepth: Universal monocular metric depth estimation," 2024. [Online]. Available: <https://arxiv.org/abs/2403.18913>
- [19] S. F. Bhat, R. Birkel, D. Wofk, P. Wonka, and M. Müller, "Zoedepth: Zero-shot transfer by combining relative and metric depth," 2023. [Online]. Available: <https://arxiv.org/abs/2302.12288>
- [20] W. Yin, C. Zhang, H. Chen, Z. Cai, G. Yu, K. Wang, X. Chen, and C. Shen, "Metric3d: Towards zero-shot metric 3d prediction from a single image," 2023. [Online]. Available: <https://arxiv.org/abs/2307.10984>
- [21] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, "Depth anything v2," 2024. [Online]. Available: <https://arxiv.org/abs/2406.09414>
- [22] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, "Depth anything: Unleashing the power of large-scale unlabeled data," 2024. [Online]. Available: <https://arxiv.org/abs/2401.10891>
- [23] V. Guizilini, I. Vasiljevic, D. Chen, R. Ambrus, and A. Gaidon, "Towards zero-shot scale-aware monocular depth estimation," 2023. [Online]. Available: <https://arxiv.org/abs/2306.17253>
- [24] B. Ke, A. Obukhov, S. Huang, N. Metzger, R. C. Daudt, and K. Schindler, "Repurposing diffusion-based image generators for monocular depth estimation," 2024. [Online]. Available: <https://arxiv.org/abs/2312.02145>
- [25] M. Gui, J. S. Fischer, U. Prestel, P. Ma, D. Kotovenko, O. Grebenkova, S. A. Baumann, V. T. Hu, and B. Ommer, "Depthfm: Fast monocular depth estimation with flow matching," 2024. [Online]. Available: <https://arxiv.org/abs/2403.13788>
- [26] S. Farooq Bhat, I. Alhashim, and P. Wonka, "Adabins: Depth estimation using adaptive bins," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Jun. 2021. [Online]. Available: <http://dx.doi.org/10.1109/CVPR46437.2021.00400>
- [27] S. F. Bhat, I. Alhashim, and P. Wonka, "Localbins: Improving depth estimation by learning local distributions," 2022. [Online]. Available: <https://arxiv.org/abs/2203.15132>
- [28] A. Agarwal and C. Arora, "Attention attention everywhere: Monocular depth prediction with skip attention," 2022. [Online]. Available: <https://arxiv.org/abs/2210.09071>
- [29] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," 2021. [Online]. Available: <https://arxiv.org/abs/2103.13413>
- [30] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from rgbd images," in *Computer Vision – ECCV 2012*, A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 746–760.
- [31] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 3354–3361.
- [32] V. Guizilini, R. Ambrus, W. Burgard, and A. Gaidon, "Sparse auxiliary networks for unified monocular depth prediction and completion," 2021. [Online]. Available: <https://arxiv.org/abs/2103.16690>
- [33] H. Yin, T. Chen, L. Qu, and B. Cui, "On-device recommender systems: A tutorial on the new-generation recommendation paradigm," in *Companion Proceedings of the ACM Web Conference 2024*, ser. WWW '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 1280–1283. [Online]. Available: <https://doi.org/10.1145/3589335.3641250>
- [34] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," 2020. [Online]. Available: <https://arxiv.org/abs/1907.01341>
- [35] W. Yin, Y. Liu, and C. Shen, "Virtual normal: Enforcing geometric constraints for accurate and robust depth prediction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 7282–7295, 2022.
- [36] A. Eftekhar, A. Sax, R. Bachmann, J. Malik, and A. Zamir, "OmniData: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans," 2021. [Online]. Available: <https://arxiv.org/abs/2110.04994>
- [37] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," 2020. [Online]. Available: <https://arxiv.org/abs/1911.05722>
- [38] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," 2020. [Online]. Available: <https://arxiv.org/abs/2002.05709>
- [39] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," 2021. [Online]. Available: <https://arxiv.org/abs/2103.00020>

- [40] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, "Dinov2: Learning robust visual features without supervision," 2024. [Online]. Available: <https://arxiv.org/abs/2304.07193>
- [41] J. Lu, F. Richter, and M. C. Yip, "Markerless camera-to-robot pose estimation via self-supervised sim-to-real transfer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 296–21 306.
- [42] N. Karaev, I. Rocco, B. Graham, N. Neverova, A. Vedaldi, and C. Rupprecht, "Cotracker: It is better to track together," 2023. [Online]. Available: <https://arxiv.org/abs/2307.07635>
- [43] J. Lu, F. Richter, and M. C. Yip, "Pose estimation for robot manipulators via keypoint optimization and sim-to-real transfer," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4622–4629, 2022.
- [44] W. Yuan, X. Gu, Z. Dai, S. Zhu, and P. Tan, "New crfs: Neural window fully-connected crfs for monocular depth estimation," 2022. [Online]. Available: <https://arxiv.org/abs/2203.01502>