Test3R: Learning to Reconstruct 3D at Test Time

Yuheng Yuan¹ Qiuhong Shen¹ Shizun Wang¹ Xingyi Yang^{2,1} Xinchao Wang^{1*}

¹National University of Singapore, ²The Hong Kong Polytechnic University yuhengyuan@u.nus.edu, xingyi.yang@polyu.edu.hk, xinchao@nus.edu.sg

*Project page: https://test3r-nop.github.io/**



Figure 1: Given a set of images of a specific scene, **Test3R** improves the quality of reconstruction by maximizing the consistency between the pointmaps generated from multiple image pairs.

Abstract

Dense matching methods like DUSt3R regress pairwise pointmaps for 3D reconstruction. However, the reliance on pairwise prediction and the limited generalization capability inherently restrict the global geometric consistency. In this work, we introduce $\mathbf{Test3R}$, a surprisingly simple test-time learning technique that significantly boosts geometric accuracy. Using image triplets (I_1, I_2, I_3) , $\mathbf{Test3R}$ generates reconstructions from pairs (I_1, I_2) and (I_1, I_3) . The core idea is to optimize the network at test time via a self-supervised objective: maximizing the geometric consistency between these two reconstructions relative to the common image I_1 . This ensures the model produces cross-pair consistent outputs, regardless of the inputs. Extensive experiments demonstrate that our technique significantly outperforms previous state-of-the-art methods on the 3D reconstruction and multiview depth estimation tasks. Moreover, it is universally applicable and nearly cost-free, making it easily applied to other models and implemented with minimal test-time training overhead and parameter footprint. Code is available at https://github.com/nopQAQ/Test3R.

^{*}Corresponding Author

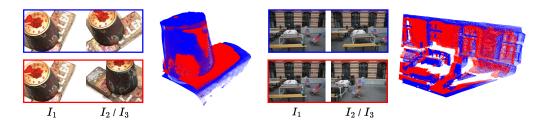


Figure 2: **Inconsistency Study.** On the left are two image pairs sharing the same reference view I_1 but with different source views I_2 and I_3 . On the right are the corresponding point maps, with each color indicating the respective image pair.

1 Introduction

3D reconstruction from multi-view images is a cornerstone task in computer vision. Traditionally, this process has been achieved by assembling classical techniques such as keypoint detection [1–3] and matching [4, 5], robust camera estimation [4, 6], Structure-from-Motion(SfM), Bundle Adjustment(BA) [7–9], and dense Multi-View Stereo [10, 11]. Although effective, these multi-stage methods require significant engineering effort to manage the entire process. This complexity inherently constrains their scalability and efficiency.

Recently, dense matching methods, such as DUSt3R [12] and MAST3R [13], have emerged as compelling alternatives. At its core, DUSt3R utilizes a deep neural network trained to *predict dense correspondences* between image pairs in an end-to-end fashion. Specifically, DUSt3R takes in two images and, for each, predicts a *pointmap*. Each pointmap represents the 3D coordinates of every pixel, as projected into a common reference view's coordinate system. Once pointmaps are generated from multiple views, DUSt3R aligns them by optimizing the registration of these 3D points. This process recovers the camera pose for each view and reconstructs the overall 3D geometry.

Despite its huge success, this *pair-wise prediction* paradigm is inherently problematic. Under such a design, the model considers only two images at a time. Such a constraint leads to several issues.

To investigate this, we compare the pointmaps of image I_1 but with different views I_2 and I_3 in Figure 2. It demonstrates that the predicted pointmaps are *imprecise* and *inconsistent*. Firstly, the precision of geometric predictions can suffer because the model is restricted to inferring scene geometry from just one image pair. This is especially true for *short-baseline* cases [14], where small camera movement leads to poor triangulation and thus inaccurate geometry. Second, reconstructing an entire scene requires pointmaps from multiple image pairs. Unfortunately, these individual pairwise predictions may not be mutually consistent. For example, the pointmap predicted from (I_1, I_2, Y_8) may not align with the prediction from (I_1, Y_2, I_3) , as highlighted by the color difference in Figure 2. This local inconsistency further leads to discrepancies in the overall reconstruction. What makes things worse, the model, like many deep learning systems, struggles to generalize to new or diverse scenes. Such limitations directly exacerbate the previously discussed problems of precision and inter-pair consistency. Consequently, even with a final global refinement stage, inaccurate pointmaps lead to persistent errors.

To address these problems, in this paper, we present **Test3R**, a novel yet strikingly simple solution for 3D reconstruction, operating entirely *at test time*. Its core idea is straightforward: Maximizing the consistency between the reconstructions generated from multiple image pairs. This principle is realized through two basic steps:

- Given image triplets (I_1, I_2, I_3) , Test3R first estimates two initial pointmaps with respect to I_1 : X_1 from pairs (I_1, I_2) and X_2 from (I_1, I_3) .
- Test3R optimizes the network, so that the two pointmaps are cross-pair consistent, i.e., $X_1 \approx X_2$. Critically, this optimization is performed at test time via prompt tuning [15].

Despite its simplicity, Test3R offers a robust solution to all challenges mentioned above. It ensures consistency by aligning local two-view predictions, which resolves inconsistencies. This same mechanism also improves geometric precision: if a pointmap from short-baseline images is imprecise,

Test3R pushes it closer to an overall global prediction, which reduces errors. Finally, Test3R adapts to new, unseen scenes, minimizing its errors on unfamiliar data.

We evaluated Test3R on the DUSt3R for 3D reconstruction and multi-view depth estimation. Test3R performs exceptionally well across diverse datasets, improving upon vanilla DUSt3R to achieve competitive or state-of-the-art results in both tasks. Surprisingly, for multi-view depth estimation, Test3R even surpasses baselines requiring camera poses and intrinsics, as well as those trained on the same domain. This further validates our model's robustness and efficacy.

The best part is that Test3R is universally applicable and nearly cost-free for these pair-wise methods. This means it can easily be applied to other models sharing a similar pipeline. We validated this by incorporating our design into two categories of models, pairwise methods like MonST3R [16] and MAST3R [13], and multi-view models like VGGT [17]. Experimental results confirmed substantial performance improvements for both models.

The contributions of this work are as follows:

- We introduce **Test3R**, a novel yet simple solution to learn the reconstruction at test time. It optimizes the model via visual prompts to maximize the cross-pair consistency. It provides a robust solution to the challenges of the pairwise prediction paradigm and limited generalization capability.
- We conducted comprehensive experiments across several downstream tasks on the DUSt3R. Experiment results demonstrate that Test3R not only improves the reconstruction performance compared to vanilla DUSt3R but also outperforms a wide range of baselines.
- Our design is universally applicable and nearly cost-free. It can be easily applied to other models and implemented with minimal test-time training overhead and parameter footprint.

2 Related Work

2.1 Multi-view Stereo

Multi-view Stereo(MVS) aims to densely reconstruct the geometry of a scene from multiple overlapping images. Traditionally, all camera parameters are often estimated with SfM [18], as the given input. Existing MVS approaches can generally be classified into three categories: traditional handcrafted [11, 19-21], global optimization [22-25], and learning-based methods [10, 26-29]. Recently, DUSt3R [12] has attracted significant attention as a representative of learning-based methods. It attempts to estimate dense pointmaps from a pair of views without any explicit knowledge of the camera parameters. Subsequent tremendous works focus on improving its efficiency [17, 30, 31], quality [13, 17, 32], and broadening its applicability to dynamic reconstruction [16, 33–35] and 3D perception [36]. The majority employ the pairwise prediction strategy introduced by DUSt3R [12]. However, the pair-wise prediction paradigm is inherently problematic. It leads to low precision and mutually inconsistent pointmaps. Furthermore, the limited generalization capability of the model exacerbates these issues. This challenge continues even with the latest models [17, 37], which can process multiple images in a single forward pass. While potentially more robust, these newer approaches demand significantly larger resources for training and, importantly, still face challenges in generalizing to unseen environments. To this end, we introduce a novel test-time training technique. This simple design ensures the cross-pairs consistency by aligning local two-view predictions to push the pointmaps closer to an overall global prediction, which addresses all challenges mentioned above.

2.2 Test-time Training

The idea of training on unlabeled test data dates back to the 1990s [38], called transductive learning. As Vladimir Vapnik [39] famously stated, "Try to get the answer that you really need but not a more general one", this principle has been widely applied to SVMs [40, 41] and recently in large language models [42]. Another early line of work is local learning [43, 44]: for each test input, a "local" model is trained on the nearest neighbors before a prediction is made. Recently, Test-time training(TTT) [45] proposes a general framework for test-time training with self-supervised learning, which produces a different model for every single test input through the self-supervision task. This strategy allows the model trained on the large-scale datasets to adapt to the target domain at test time. Many other works have followed this framework since then [46–49]. Inspired by these studies, we introduce Test3R, a novel yet simple technique that extends the test-time training paradigm to the 3D reconstruction

domain. Our model exploits the cross-pairs consistency as a strong self-supervised objective to optimize the model parameters at test time, thereby improving the final quality of reconstruction.

2.3 Prompt tuning

Prompt tuning was first proposed as a technique that appends learnable textual prompts to the input sequence, allowing pre-trained language models to adapt to downstream tasks without modifying the backbone parameters [50]. In follow-up research, a portion of studies [51, 52] explored strategies for crafting more effective prompt texts, whereas others [53–55] proposed treating prompts as learnable, task-specific continuous embeddings, which are optimized via gradient descent during fine-tuning referred to as Prompt Tuning. In recent years, prompt tuning has also received considerable attention in the 2D vision domain. Among these, Visual Prompt Tuning (VPT) [15] has gained significant attention as an efficient approach specifically tailored for vision tasks. It introduces a set of learnable prompt tokens into the pretrained model and optimizes them using the downstream task's supervision while keeping the backbone frozen. This strategy enables the model to transfer effectively to downstream tasks. In our study, we leverage the efficient fine-tuning capability of VPT to optimize the model to ensure the pointmaps are cross-view consistent. This design makes our model nearly cost-free, requiring minimal test-time training overhead and a small parameter footprint.

3 Preliminary of DUSt3R

Given a set of images $\{I_k\}$ of a specific scene, DUSt3R [12] achieves high precision 3D reconstruction by predicting pairwise pointmaps of all views and global alignment.

Pairwise prediction. Briefly, DUSt3R takes a pair of images, $I^1, I^2 \in \mathbb{R}^{W \times H \times 3}$ as input and outputs the corresponding pointmaps $X^{1,1}, X^{2,1} \in \mathbb{R}^{W \times H \times 3}$ which are expressed in the same coordinate frame of I^1 . In our paper, we refer to the viewpoint of I^1 as the reference view, while the other is the source view. Therefore, the pointmaps $X^{1,1}, X^{2,1}$ can be denoted as $X^{ref,ref}, X^{src,ref}$, respectively.

In more detail, these two input images I^{ref} , I^{src} are first encoded by the same weight-sharing ViT-based model [56] with N_e layers to yield two token representations F^{ref} and F^{src} :

$$F^{ref} = Encoder(I^{ref}), \quad F^{src} = Encoder(I^{src})$$
 (1)

After encoding, the network reasons over both of them jointly in the decoder. Each decoder block also attends to tokens from the other branch:

$$G_{i}^{ref} = DecoderBlock_{i}^{ref}(G_{i-1}^{ref}, G_{i-1}^{src}) \tag{2} \label{eq:2}$$

$$G_i^{src} = DecoderBlock_i^{src}(G_{i-1}^{src}, G_{i-1}^{ref})$$

$$\tag{3}$$

where $i=1,\cdots,N_d$ for a decoder with N_d decoder layers and initialized with encoder tokens $G_0^{ref}=F^{ref}$ and $G_0^{src}=F^{src}$. Finally, in each branch, a separate regression head takes the set of decoder tokens and outputs a pointmap and an associated confidence map:

$$X^{ref,ref}, C^{ref,ref} = Head^{ref}(G_0^{ref}, \dots, G_{N_d}^{ref}), \tag{4}$$

$$X^{src,ref}, C^{src,ref} = Head^{src}(G_0^{src}, \dots, G_{N_d}^{src}).$$

$$\tag{5}$$

Global alignment. After predicting all the pairwise pointmaps, DUSt3R introduces a global alignment to handle pointmaps predicted from multiple images. For the given image set $\{\mathbf{I}_t^i\}_{i=1}^{N_t}$, DUSt3R first constructs a connectivity graph $\mathcal{G}(\mathcal{V},\mathcal{E})$ for selecting pairwise images, where the vertices \mathcal{V} represent N_t images and each edge $e \in \mathcal{E}$ is an image pair. Then, it estimates the depth maps $D := \{\mathbf{D}_k\}$ and camera pose $\pi := \{\pi_k\}$ by

$$\underset{\mathbf{D},\pi,\sigma}{\operatorname{arg\,min}} \sum_{e\in\mathcal{E}} \sum_{v\in e} \mathbf{C}_{v}^{e} \left\| \mathbf{D}_{v} - \sigma_{e} P_{e}(\pi_{v}, \mathbf{X}_{v}^{e}) \right\|_{2}^{2}, \tag{6}$$

where $\sigma = \{\sigma_e\}$ are the scale factors defined on the edges, $P_e(\pi_v, \mathbf{X}_v^e)$ means projecting the predicted pointmap \mathbf{X}_v^e to view v using poses π_v to get a depth map. The objective function in eq. (6) explicitly constrains the geometry alignment between frame pairs, aiming to preserve cross-view consistency in the depth maps.

4 Methods

Test3R is a test-time training technique that adapts DUSt3R [12] to challenging test scenes. It improves reconstruction by maximizing cross-pair consistency. We begin by analyzing the root cause of inconsistency in Sec. 4.1. In Sec. 4.2, we establish the core problem and define the test-time training objective. Finally, we employ prompt tuning for efficient test-time adaptation in Sec. 4.3.

4.1 Cross-pair Inconsistency

DUSt3R [12] aims to achieve consistency through global alignment; however, the inaccurate and inconsistent pointmaps lead to persistent errors, significantly compromising the effectiveness of global alignment.

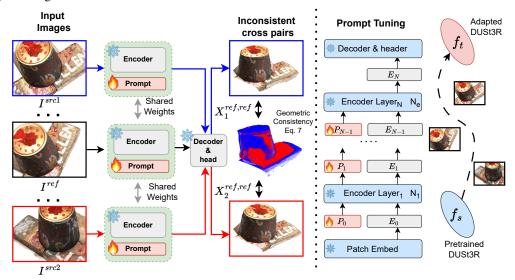


Figure 3: **Overview of Test3R.** The primary goal of Test3R is to adapt a pretrained reconstruction model f_s to the specific distribution of test scenes f_t . It achieves this goal by optimizing a set of visual prompts at test time through a self-supervised training objective that maximizes cross-pair consistency between $X_1^{ref,ref}$ and $X_2^{ref,ref}$.

Therefore, we show a qualitative analysis of the quality of pointmaps on the DTU [57] and ETH3D [58] datasets. Specifically, we compare the pointmap for the same reference view but paired with two different source views, and align these two pointmaps to the same coordinate system using Iterative Closest Point (ICP). The result is shown in Figure 2. On the left are two image pairs sharing the same reference view but with different source views. On the right are the corresponding pointmaps, with each color indicating the respective image pair.

<u>Observations</u>. These two predicted pointmaps of the reference view exhibit inconsistencies, as highlighted by the presence of large regions with inconsistent colors in 3D space. Ideally, if these pointmaps are consistent, they should be accurate enough to align perfectly in 3D space, resulting in a single, unified color (either blue or red). This result indicates that DUSt3R may produce different pointmaps for the same reference view when paired with different source views.

In our view, this phenomenon stems from the problematic pair-wise prediction paradigm. First, since only two views are provided as input at each prediction step, the scene geometry is estimated solely based on visual correspondences between a single image pair. Therefore, the model produces inaccurate pointmaps. Second, all predicted pointmaps are mutually inconsistent individual pairs. For different image pairs, their visual correspondences are also different. As a result, DUSt3R may produce inconsistent pointmaps for the same reference view when paired with different source views due to the different correspondences. This issue significantly hinders the effectiveness of subsequent global alignment and further leads to discrepancies in the overall reconstruction. What's worse, the limited generalization capability of DUSt3R further exacerbates the above issues of low precision and cross-pair inconsistency.

4.2 Triplet Objective Made Consistent

The inconsistencies observed above highlight a core limitation of the pairwise prediction paradigm. Specifically, DUSt3R may produce different pointmaps for the same reference view when paired with different source views. This motivates a simple but effective idea: enforce triplet consistency across these pointmaps directly at test time, as shown in Figure 3.

Definition. We first describe the definition of test-time training on the 3D reconstruction task, where only images $\{I_t^i\}_{i=1}^{N_t}$ from the test scene are available. During training time training phase, N_s labeled samples $\{I_s^i, \bar{X}_s^i\}_{i=1}^{N_s}$ collected from various scenes are given, where $I_s^i \in \mathcal{I}_s$ and $\bar{X}_s^i \in \bar{\mathcal{X}}_s$ are images and the corresponding pointmaps derived from the ground-truth depth $\bar{D}_s \in \bar{\mathcal{D}}_s$. Furthermore, we denote DUSt3R [12], parameterized by θ , as the model trained to learn the reconstruction function $f_s:\mathcal{I}_s\to\bar{\mathcal{X}}_s$. Subsequently, during test time training phase, only unlabeled images $\{I_t^i\}_{i=1}^{N_t}$ from test scene are available, where $I_t^i\in\mathcal{I}_t$. Our goal is to optimize the model f_s to the specific scene $f_t:\mathcal{I}_t\to\bar{\mathcal{X}}_t$ at test time. This is achieved by minimizing the self-supervised training objective ℓ .

Specifically, our core training objective is to maximize the geometric consistency by aligning the pointmaps of the reference view when paired with different source views. For a set of images $\{I_t^i\}_{i=1}^{N_t}$ from the specific scene, we consider a triplet consisting of one reference view and two different source views, denoted as $(I^{ref},I^{src1},I^{src2})$. Subsequently, Test3R forms two reference—source view pairs (I^{ref},I^{src1}) and (I^{ref},I^{src2}) from this triplets. These reference—source view pairs are then fed into the Test3R independently to predict pointmaps of reference views under different source view conditions in the same coordinate frame of I^{ref} , denoted as $\mathbf{X}_1^{ref,ref}$ and $\mathbf{X}_2^{ref,ref}$. Finally, we construct the training objective by aligning these two inconsistent pointmaps, formulated as:

$$\ell = \left\| X_1^{ref,ref} - X_2^{ref,ref} \right\|_1. \tag{7}$$

With this objective, we can collectively compose triplets from a large number of views of an unseen 3D scene at test time. It guides the model to successfully resolve the limitations mentioned in Section 4.1. For inconsistencies, it ensures consistency by aligning the local two-view predictions. Meanwhile, it also pushes the predicted pointmap closer to an overall global prediction to mitigate the inaccuracy. Moreover, by optimizing for the specific scene at test time, it enables the model to adapt to the distribution of that scene.

4.3 Visual Prompt Tuning for Test Time Training

After the self-supervised training objective is defined, effectively modulating the model during test-time training for specific scenes remains a non-trivial challenge. During the test-time training phase, it only relies on unsupervised training objectives. However, these objectives are often noisy and unreliable, which makes the model prone to overfitting and may lead to training collapse, especially when only a limited number of images are available for the current scene. Fortunately, similar issues has been partially explored in the 2D vision community. In these works, visual prompt tuning [15] has demonstrated strong effectiveness in domain adaptation in 2D classification tasks [59]. It utilizes a set of learnable continuous parameters to learn the specific knowledge while retaining the knowledge learned from large-scale pretraining. Motivated by this, we explore the use of visual prompts as a carrier to learn the geometric consistency for specific scenes.

Specifically, we incorporate a set of learnable prompts into the encoder of DUSt3R [12]. Consider an encoder of DUSt3R with N_e standard Vision Transformer(ViT) [56] layers, an input image is first divided into fixed-sized patches and then embedded into d-dimensional tokens $\mathbf{E_0} = \{\mathbf{e}_0^k \in \mathbb{R}^D | k \in \mathbb{N}, 1 \leq k \leq N_t\}$, where N_t is the length of image patch tokens. Subsequently, to optimize the model, we introduce a set of learnable prompt tokens $\{\mathbf{P}_{i-1}\}_{i=1}^{N_e}$ into each Transformer layer. For i-th transformer layer, the prompt tokens are denoted as $\mathbf{P}_{i-1} = \{\mathbf{p}_{i-1}^k \in \mathbb{R}^D | k \in \mathbb{N}, 1 \leq k \leq N_p\}$, where N_p is the length of prompt tokens. Therefore, the encoder layer augmented by visual prompts is formulated as:

$$[\underline{\mathbf{E}}_{\mathbf{i}}] = L_i([\mathbf{P}_{\mathbf{i}-\mathbf{1}}, \mathbf{E}_{\mathbf{i}-\mathbf{1}}]) \tag{8}$$

where P_{i-1} and E_{i-1} are learnable prompt tokens and image patch tokens at i-1-th Transformer layer.

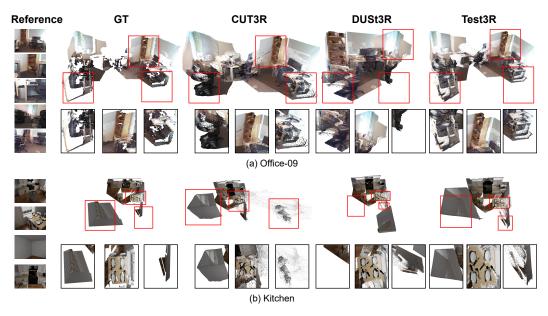


Figure 4: Qualitative Comparison on 3D Reconstruction.

Test-time training. We only fine-tune the parameters of the prompts, while all other parameters are fixed. This strategy enables our model to maximize geometric consistency by optimizing the prompts at test time while retaining the reconstruction knowledge acquired from large-scale datasets training within the unchanged backbone.

5 Experiment

We evaluate our method across a range of 3D tasks, including 3D Reconstruction (Section 5.1) and Multi-view Depth (Section 5.2). Moreover, we discuss the generalization of Test3R and the prompt design (Section 5.3). Additional experiments and detailed information are provided in the appendix.

Baselines. Our primary baseline is DUSt3R [12], which serves as the backbone of our technique in the experiment. Subsequently, we select different baselines for the specific tasks to comprehensively evaluate the performance of our proposed method. For the 3D reconstruction task, which is the primary focus of the majority of 3R-series models, we compared our method with current mainstream approaches to evaluate its effectiveness. It includes MAST3R [13], MonST3R [16], CUT3R [35] and Spann3R [31]. All of these models are follow-up works building on the foundation established by DUSt3R [12]. Furthermore, for the multi-view depth estimation task, we not only compare our model with baselines [60, 61] that do not require camera parameters but also evaluate our model against methods [9, 11, 28, 60–63, 63, 64] that rely on camera parameters or trained on datasets from the same distribution to demonstrate the effectiveness of our technique.

5.1 3D Reconstruction

We utilize two scene-level datasets, 7Scenes [65] and NRGBD [66] datasets. We follow the experiment setting on the CUT3R [35], and employ several commonly used metrics: Accuracy (Acc), Completion (Comp), and Normal Consistency (NC) metrics. Each scene has only 3 to 5 views available for the 7Scenes [65] dataset and 2 to 4 views for NRGBD [66] dataset. This is a highly challenging experimental setup, as the overlap between images in each scene is minimal, demanding a strong scene reconstruction capability.

Quantitative Results. The quantitative evaluation is shown in Table 1. Compared to vanilla DUSt3R [12], our model demonstrates superior performance, outperforming DUSt3R on the majority of evaluation metrics, particularly in terms of mean accuracy and completion. Moreover, our approach achieves comparable or even superior results compared to mainstream methods. Only CUT3R [35] and MAST3R [13] outperform our approach on several metrics. This demonstrates the effectiveness of our test-time training strategy.

Qualitative Results. The qualitative results are shown in Figure 4. We compare our method with CUT3R [35] and DUSt3R [12] on the *Office* and *Kitchen* scenes from the 7Scenes [65] and NRGBD [66] datasets, respectively. We observe that DUSt3R incorrectly regresses the positions of scene views, leading to errors in the final scene reconstruction. In contrast, our model achieves more reliable scene reconstructions. This improvement is particularly evident in the statue in the *Office* scene and the wall in the *Kitchen* scene. For these two objects, the reconstruction results from DUSt3R are drastically different from the ground truth. Compared to CUT3R [35], the current state-of-the-art in 3D reconstruction, we achieve better reconstruction results. Specifically, we effectively avoid the generation of outliers, resulting in more accurate pointmaps. Details can be seen in the red bounding boxes as shown in Figure 4.

Table 1: 3D reconstruction comparison on 7Scenes and NRGBD datasets.

	7Scenes						NRGBD					
	Acc↓		Comp↓		NC↑		Acc↓		Comp↓		NC↑	
Method	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.
MAST3R [13]	0.189	0.109	0.211	0.110	0.687	0.766	0.085	0.033	0.063	0.028	0.794	0.928
MonST3R [16]	0.240	0.180	0.268	0.167	0.672	0.758	0.272	0.114	0.287	0.110	0.758	0.843
Spann3R [31]	0.298	0.226	0.205	0.112	0.650	0.730	0.416	0.323	0.417	0.285	0.684	0.789
CUT3R [35]	0.126	0.047	0.154	0.031	0.727	0.834	0.099	0.031	0.076	0.026	0.837	0.971
DUSt3R [12]	0.146	0.078	0.181	0.067	0.736	0.839	0.144	0.019	0.154	0.018	0.871	0.982
Test3R(Ours)	0.105	0.051	0.136	0.035	0.746	0.855	0.083	0.021	0.079	0.019	0.870	0.983

Table 2: Multi-view depth evaluation. (Parentheses) denote training on data from the same domain.

Method	GT Pose	GT GT Range Intrinsics		Align	DTU rel \downarrow $\tau \uparrow$		ETH3D rel $\downarrow \tau \uparrow$		AVG rel $\downarrow \tau \uparrow$	
COLMAP [9, 11]		×		×	0.7	96.5	16.4	55.1	8.6	75.8
COLMAP Dense [9, 11]	./	×	<i>,</i>	×	20.8	69.3	89.8	23.2	55.3	46.3
MVSNet [28]	1	7	· /	×	(1.8)	(86.0)	35.4	31.4	18.6	58.7
Vis-MVSSNet [63]	· /	· /	· /	×	(1.8)	(87.4)	10.8	43.3	6.3	65.4
MVS2D ScanNet [64]	· /	· /	· /	X	17.2	9.8	27.4	4.8	22.3	7.3
MVS2D DTU [64]	<i>\</i>	· /	<i>'</i>	×	(3.6)	(64.2)	99.0	11.6	51.3	37.9
DeMoN [60]	✓	×	✓	×	23.7	11.5	19.0	16.2	21.4	13.9
DeepV2D KITTI [61]	✓	×	✓	×	24.6	8.2	30.1	9.4	27.4	8.8
DeepV2D ScanNet [61]	✓	×	✓	×	9.2	27.4	18.7	28.7	14.0	28.1
MVS2D ScanNet [64]	✓	×	✓	×	5.0	57.9	30.7	14.4	17.9	36.2
Robust MVD Baseline [62]	\checkmark	×	✓	×	2.7	82.0	9.0	42.6	5.9	62.3
DeMoN [60]	×	×	√	t	21.8	16.6	17.4	15.4	19.6	16.0
DeepV2D KITTI [61]	×	×	✓	med	24.8	8.1	27.1	10.1	26.0	9.1
DeepV2D ScanNet [61]	×	×	✓	med	7.7	33.0	11.8	29.3	9.8	62.3
DUSt3R [1]	×	×	×	med	3.3	69.9	3.3	73.0	3.3	71.5
Test3R(Ours)	×	×	×	med	2.0	84.1	3.2	74.0	2.6	79.1

5.2 Multi-view Depth

Following RobustMVD [62], performances are measured on the object-centric dataset DTU [57] and scene-centric dataset ETH3D [58]. To evaluate the depth map, we report the Absolute Relative Error (rel) and the Inlier Ratio (τ) at a threshold of 3% on each test set and the averages across all test sets.

Quantitative Results. The quantitative evaluation is shown in Table 2. On the DTU dataset, our model significantly improves upon the performance of vanilla DUSt3R, reducing the Absolute Relative Error by 1.3 and increasing the Inlier Ratio by 14.2. Similarly, on the ETH3D dataset, our model also demonstrates comparable improvements, achieving state-of-the-art performance on this challenging benchmark as well. Notably, our model surpasses the majority of methods that rely on camera poses and intrinsic parameters, and the models trained on the dataset from the same domain. This indicates that our approach effectively captures scene-specific global information and enables the adaptation of the distribution of test scenes, thereby significantly improving the quality of the depth maps.

Qualitative Results. The qualitative result is shown in Figure 5. We present the depth map on the key view, following RobustMVD [62]. We observe that Test3R effectively improves the accuracy of depth estimation compared to DUSt3R and RobustMVD [62] with camera parameters. Specifically, Test3R captures more fine-grained details, including the computer chassis and table. Additionally, on the white-background DTU dataset, Test3R effectively understands scene context, allowing it to accurately estimate the depth of background regions.

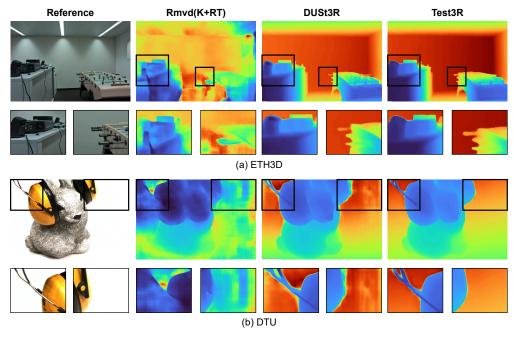


Figure 5: Qualitative Comparison on Multi-view Depth.

5.3 Ablation Study and Analysis

5.3.1 Framework Generalization.

To demonstrate the generalization ability of our proposed technique, we applied Test3R to MAST3R [13] and MonST3R [16], and evaluated the performances on the 7Scenes [65] dataset. As shown in Table 3, Test3R effectively improves the performance of MAST3R and MonST3R on 3D reconstruction task. This demonstrates the generalization ability of our technique, which can be applied to other models sharing a similar pipeline.

5.3.2 Ablation on Visual Prompt.

We introduce a model variant, Test3R-S, and conduct an ablation study to evaluate the impact of visual prompts. For Test3R-S, the prompts are only inserted into the first Transformer layer, accompany the image tokens through the encoding process, and are then discarded.

Table 3: Generalization Study.

	7Scenes						
	Ac	cc.	Comp↓		N	C↑	
Method	Mean	Med.	Mean	Med.	Mean	Med.	
MAST3R [13] MAST3R(w. Test3R)	0.189 0.179	0.109 0.108	0.211 0.177	0.110 0.059	0.687 0.702	0.766 0.788	
MonST3R [16] MonST3R(w. Test3R)	0.240 0.218	0.180 0.167	0.268 0.251	0.167 0.160	0.672 0.687	0.758 0.775	

Table 4: **Ablation study on Visual Prompt.**

	Prompts Length									
Varients	8		16		:	32	64			
	Acc↓	Comp↓	Acc	$\text{Comp}{\downarrow}$	Acc↓	Comp↓	Acc↓	Comp↓		
Test3R-S	0.133	0.142	0.125	0.159	0.120	0.158	0.119	0.163		
Test3R	0.118	0.131	0.122	0.155	0.105	0.136	0.149	0.170		

The result is shown in Table 4. Both Test3R-S and Test3R effectively improve model performance, compared to vanilla DUSt3R. For prompt length, we observe that when the number of prompts is small, increasing the prompt length can enhance the ability of Test3R to improve reconstruction quality. However, as the prompt length increases, the number of trainable parameters also grows, making it more challenging to converge within the same number of iterations, thereby reducing their overall effectiveness. For prompt insertion depth, we observe that Test3R, which uses distinct prompts at each layer, demonstrates superior performance. This is because the feature distributions vary across each layer of the encoder of DUSt3R, making layer-specific prompts more effective for fine-tuning. However, as the number of prompt parameters increases, Test3R becomes more susceptible to optimization challenges compared to Test3R-S, leading to a faster performance decline.

6 Conclusion

In this paper, we present Test3R, a novel yet strikingly simple solution that learns to reconstruct at test time. It maximizes the cross-pair consistency via optimizing a set of visual prompts at test time. This design successfully mitigates the reconstruction quality degradation caused by the pairwise predictions paradigm and limited generalization capability. Extensive experiments show that our simple design not only effectively improves model performance but also achieves state-of-the-art performance across various tasks. Moreover, our technique is universally applicable and nearly cost-free, which can be widely applied to different models and implemented with minimal test-time training overhead and parameter footprint.

Acknowledgement

This project is supported by the National Research Foundation, Singapore, under its Medium Sized Center for Advanced Robotics Technology Innovation.

References

- [1] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint description and detection of local features. In *Proceedings of the ieee/cyf conference on computer vision and pattern recognition*, pages 8092–8101, 2019.
- [2] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018.
- [3] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004.
- [4] Eric Brachmann and Carsten Rother. Neural-guided ransac: Learning where to sample model hypotheses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4322–4331, 2019.
- [5] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. Lightglue: Local feature matching at light speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17627–17638, 2023.
- [6] Chen Zhao, Yixiao Ge, Feng Zhu, Rui Zhao, Hongsheng Li, and Mathieu Salzmann. Progressive correspondence pruning by consensus learning. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pages 6464–6473, 2021.
- [7] David J Crandall, Andrew Owens, Noah Snavely, and Daniel P Huttenlocher. Sfm with mrfs: Discrete-continuous optimization for large-scale structure from motion. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2841–2853, 2012.
- [8] Philipp Lindenberger, Paul-Edouard Sarlin, Viktor Larsson, and Marc Pollefeys. Pixel-perfect structure-from-motion with featuremetric refinement. In Proceedings of the IEEE/CVF international conference on computer vision, pages 5987–5997, 2021.
- [9] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016.
- [10] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2495–2504, 2020.
- [11] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III 14*, pages 501–518. Springer, 2016.
- [12] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024.

- [13] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer, 2024.
- [14] Masatoshi Okutomi and Takeo Kanade. A multiple-baseline stereo. IEEE Transactions on pattern analysis and machine intelligence, 15(4):353–363, 1993.
- [15] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European conference on computer vision*, pages 709–727. Springer, 2022.
- [16] Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. arXiv preprint arXiv:2410.03825, 2024.
- [17] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. *arXiv* preprint arXiv:2503.11651, 2025.
- [18] Shimon Ullman. The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 203(1153):405–426, 1979.
- [19] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and trends*® *in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [20] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE international conference on computer vision*, pages 873–881, 2015.
- [21] Yuesong Wang, Zhaojie Zeng, Tao Guan, Wei Yang, Zhuo Chen, Wenkai Liu, Luoyuan Xu, and Yawei Luo. Adaptive patch deformation for textureless-resilient multi-view stereo. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1621–1630, 2023.
- [22] Qiancheng Fu, Qingshan Xu, Yew Soon Ong, and Wenbing Tao. Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. Advances in Neural Information Processing Systems, 35:3403–3416, 2022.
- [23] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, pages 3504–3515, 2020.
- [24] Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5610–5619, 2021.
- [25] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. Advances in Neural Information Processing Systems, 33:2492–2502, 2020.
- [26] Zeyu Ma, Zachary Teed, and Jia Deng. Multiview stereo with cascaded epipolar raft. In *European Conference on Computer Vision*, pages 734–750. Springer, 2022.
- [27] Rui Peng, Rongjie Wang, Zhenyu Wang, Yawen Lai, and Ronggang Wang. Rethinking depth estimation for multi-view stereo: A unified representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8645–8654, 2022.
- [28] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, pages 767–783, 2018.
- [29] Zhe Zhang, Rui Peng, Yuxi Hu, and Ronggang Wang. Geomvsnet: Learning multi-view stereo with geometry perception. In *Proceedings of the IEEE/CVF conference on computer vision and pattern* recognition, pages 21508–21518, 2023.
- [30] Yuzheng Liu, Siyan Dong, Shuzhe Wang, Yingda Yin, Yanchao Yang, Qingnan Fan, and Baoquan Chen. Slam3r: Real-time dense scene reconstruction from monocular rgb videos. *arXiv preprint* arXiv:2412.09401, 2024.
- [31] Hengyi Wang and Lourdes Agapito. 3d reconstruction with spatial memory. *arXiv preprint* arXiv:2408.16061, 2024.

- [32] Wonbong Jang, Philippe Weinzaepfel, Vincent Leroy, Lourdes Agapito, and Jerome Revaud. Pow3r: Empowering unconstrained 3d reconstruction with camera and scene priors. *arXiv preprint arXiv:2503.17316*, 2025.
- [33] Jiahao Lu, Tianyu Huang, Peng Li, Zhiyang Dou, Cheng Lin, Zhiming Cui, Zhen Dong, Sai-Kit Yeung, Wenping Wang, and Yuan Liu. Align3r: Aligned monocular depth estimation for dynamic videos. arXiv preprint arXiv:2412.03079, 2024.
- [34] Xingyu Chen, Yue Chen, Yuliang Xiu, Andreas Geiger, and Anpei Chen. Easi3r: Estimating disentangled motion from dust3r without training. arXiv preprint arXiv:2503.24391, 2025.
- [35] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. *arXiv preprint arXiv:2501.12387*, 2025.
- [36] Jie Hu, Shizun Wang, and Xinchao Wang. Pe3r: Perception-efficient 3d reconstruction. arXiv preprint arXiv:2503.07507, 2025.
- [37] Jianing Yang, Alexander Sax, Kevin J Liang, Mikael Henaff, Hao Tang, Ang Cao, Joyce Chai, Franziska Meier, and Matt Feiszli. Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass. arXiv preprint arXiv:2501.13928, 2025.
- [38] Alex Gammerman, Volodya Vovk, and Vladimir Vapnik. Learning by transduction. *arXiv preprint arXiv:1301.7375*, 2013.
- [39] Vladimir Vapnik. Estimation of dependences based on empirical data. Springer Science & Business Media, 2006.
- [40] Ronan Collobert, Fabian Sinz, Jason Weston, Léon Bottou, and Thorsten Joachims. Large scale transductive svms. *Journal of Machine Learning Research*, 7(8), 2006.
- [41] Thorsten Joachims. Learning to classify text using support vector machines, volume 668. Springer Science & Business Media, 2002.
- [42] Moritz Hardt and Yu Sun. Test-time training on nearest neighbors for large language models. *arXiv* preprint arXiv:2305.18466, 2023.
- [43] Léon Bottou and Vladimir Vapnik. Local learning algorithms. Neural computation, 4(6):888–900, 1992.
- [44] Hao Zhang, Alexander C Berg, Michael Maire, and Jitendra Malik. Svm-knn: Discriminative nearest neighbor classification for visual category recognition. In 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), volume 2, pages 2126–2136. IEEE, 2006.
- [45] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pages 9229–9248. PMLR, 2020.
- [46] Nicklas Hansen, Rishabh Jangir, Yu Sun, Guillem Alenyà, Pieter Abbeel, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. Self-supervised policy adaptation during deployment. arXiv preprint arXiv:2007.04309, 2020.
- [47] Yu Sun, Wyatt L Ubellacker, Wen-Loong Ma, Xiang Zhang, Changhao Wang, Noel V Csomay-Shanklin, Masayoshi Tomizuka, Koushil Sreenath, and Aaron D Ames. Online learning of unknown dynamics for model-based controllers in legged locomotion. *IEEE Robotics and Automation Letters*, 6(4):8442–8449, 2021.
- [48] Yuejiang Liu, Parth Kothari, Bastien Van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. Ttt++: When does self-supervised test-time training fail or thrive? Advances in Neural Information Processing Systems, 34:21808–21820, 2021.
- [49] Longhui Yuan, Binhui Xie, and Shuang Li. Robust test-time adaptation in dynamic scenarios. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15922– 15932, 2023.
- [50] Ben Mann, N Ryder, M Subbiah, J Kaplan, P Dhariwal, A Neelakantan, P Shyam, G Sastry, A Askell, S Agarwal, et al. Language models are few-shot learners. arXiv preprint arXiv:2005.14165, 1:3, 2020.
- [51] Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020.

- [52] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. arXiv preprint arXiv:2010.15980, 2020.
- [53] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. arXiv preprint arXiv:2104.08691, 2021.
- [54] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. arXiv preprint arXiv:2101.00190, 2021.
- [55] Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*, 2021.
- [56] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [57] Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjorholm Dahl. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision*, 120:153–168, 2016.
- [58] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multicamera videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3260–3269, 2017.
- [59] Yunhe Gao, Xingjian Shi, Yi Zhu, Hao Wang, Zhiqiang Tang, Xiong Zhou, Mu Li, and Dimitris N Metaxas. Visual prompt tuning for test-time domain adaptation. *arXiv* preprint arXiv:2210.04831, 2022.
- [60] Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. Demon: Depth and motion network for learning monocular stereo. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5038–5047, 2017.
- [61] Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. *arXiv* preprint arXiv:1812.04605, 2018.
- [62] Philipp Schröppel, Jan Bechtold, Artemij Amiranashvili, and Thomas Brox. A benchmark and a baseline for robust multi-view depth estimation. In 2022 International Conference on 3D Vision (3DV), pages 637–645. IEEE, 2022.
- [63] Jingyang Zhang, Shiwei Li, Zixin Luo, Tian Fang, and Yao Yao. Vis-mvsnet: Visibility-aware multi-view stereo network. *International Journal of Computer Vision*, 131(1):199–214, 2023.
- [64] Zhenpei Yang, Zhile Ren, Qi Shan, and Qixing Huang. Mvs2d: Efficient multi-view stereo via attentiondriven 2d convolutions. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 8574–8584, 2022.
- [65] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2930–2937, 2013.
- [66] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural rgb-d surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6290–6301, 2022.
- [67] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [68] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
- [69] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In European conference on computer vision, pages 611–625. Springer, 2012.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Test3R significantly improves the performance of DUSt3R and can achieve performance comparable to current baselines, as demonstrated in Table 1 and Table 2. Moreover, it can also be applied to other models Table 3.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations. Please see in appendix.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not involve theoretical result.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper shows the results needed and ablation study in Section 5. All of the results are reproducible.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The code will be open-sourced upon acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The detailed experiment settings are listed in Section 5, and appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: The experiment metrics are calculated on several large-scale datasets. We follow the widely adopted benchmarks and list the essential detailed quantitative results in Section 5. None of them report error bars.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The paper provides sufficient information on the computer resources in appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The research conducted in the paper conforms to the NeurIPS Code of Ethics. Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The paper discusses potential positive societal impacts and negative societal impacts.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper does not pose such a risk.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Proper credit is given to the creators or original owners of assets used in the paper, and the license and terms of use are explicitly mentioned and respected.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- · Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- · At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing and nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Not related.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.