

Depth Anything 3: Recovering the Visual Space from Any Views

Haotong Lin*, Sili Chen*, Jun Hao Liew*, Donny Y. Chen*, Zhenyu Li, Guang Shi,
Jiashi Feng, Bingyi Kang*,[†]

ByteDance Seed


[†]Project Lead, *Equal Contribution

Abstract

We present Depth Anything 3 (DA3), a model that predicts spatially consistent geometry from an arbitrary number of visual inputs, with or without known camera poses. In pursuit of minimal modeling, DA3 yields two key insights: a single plain transformer (*e.g.*, vanilla DINO encoder) is sufficient as a backbone without architectural specialization, and a singular depth-ray prediction target obviates the need for complex multi-task learning. Through our teacher-student training paradigm, the model achieves a level of detail and generalization on par with Depth Anything 2 (DA2). We establish a new visual geometry benchmark covering camera pose estimation, any-view geometry and visual rendering. On this benchmark, DA3 sets a new state-of-the-art across all tasks, surpassing prior SOTA VGGT by an average of 35.7% in camera pose accuracy and 23.6% in geometric accuracy. Moreover, it outperforms DA2 in monocular depth estimation. All models are trained exclusively on public academic datasets.

Correspondence: Bingyi Kang

Project Page: [depth-anything-3.github.io](https://github.com/depth-anything/depth-anything-3)



¹“Depth Anything 3” marks a new generation for the series, expanding from monocular to any-view inputs, built on our conviction that depth is the cornerstone of understanding the physical world.

Contents

1	Introduction	3
2	Related Work	4
3	Depth Anything 3	6
3.1	Formulation	6
3.2	Architecture	7
3.3	Training	7
3.4	Implementation Details	8
4	Teacher-Student Learning	8
4.1	Constructing the Teacher Model	8
4.2	Teaching Depth Anything 3	10
4.3	Teaching Monocular Model	10
4.4	Teaching Metric Model	10
5	Application: Feed-Forward 3D Gaussian Splattings	11
5.1	Pose-Conditioned Feed-Forward 3DGS	11
5.2	Pose-Adaptive Feed-Forward 3DGS	11
5.3	Implementation Details	12
6	Visual Geometry Benchmark	12
6.1	Benchmark Pipeline	12
6.2	Metrics	12
6.3	Datasets	13
7	Experiments	14
7.1	Comparison with State of the Art	14
7.2	Analysis for Depth Anything 3	16
7.2.1	Sufficiency of the Depth-Ray Representation	16
7.2.2	Sufficiency of a Single Plain Transformer	17
7.2.3	Ablation and Analysis	17
7.3	Analysis for Depth-Anything-3-Monocular	18
7.3.1	Teacher Model	19
7.3.2	Student Model	19
7.4	Analysis for Depth-Anything-3-Metric	19
7.5	Analysis for Feed-forward 3DGS	20
8	Conclusion and Discussion	22
A	Data Processing	31




Figure 1 Given any number of images and optional camera poses, **Depth Anything 3** reconstructs the visual space, producing consistent depth and ray maps that can be fused into accurate point clouds, resulting in high-fidelity 3D Gaussians and geometry. It significantly outperforms VGGT in multi-view geometry and pose accuracy; with monocular inputs, it also surpasses Depth Anything 2 while matching its detail and robustness.

1 Introduction

The ability to perceive and understand 3D spatial information from visual input is a cornerstone of human spatial intelligence [4] and a critical requirement for applications like robotics and mixed reality. This fundamental capability has inspired a wide array of 3D vision tasks, including Monocular Depth Estimation [23], Structure from Motion [80], Multi-View Stereo [73] and Simultaneous Localization and Mapping [58]. Despite the strong conceptual overlap between these tasks—often differing by only a single factor, such as the number of input views—the prevailing paradigm has been to develop highly specialized models for each one. While recent efforts [91, 97] have explored unified models to address multiple tasks simultaneously, they typically suffer from key limitations: they often rely on complex, bespoke architectures, are trained via joint optimization over tasks from scratch, and consequently cannot effectively leverage large-scale pretrained models.

In this work, we step back from established 3D task definitions and return to a more fundamental goal inspired by human spatial intelligence: recovering 3D structure from arbitrary visual inputs, be it a single image, multiple views of a scene, or a video stream. Forsaking intricate architectural engineering, we pursue a minimal modeling strategy guided by two central questions. First, *is there a minimal set of prediction targets, or is joint modeling across numerous 3D tasks necessary?* Second, *can a single plain transformer*

suffice for this objective? Our work provides an affirmative answer to both. We present **Depth Anything 3**, a single transformer model trained exclusively for joint **any-view depth and pose estimation** via a specially chosen ray representation. We demonstrate that this minimal approach is sufficient to reconstruct the visual space from any number of images, with or without known camera poses.

Depth Anything 3 formulates the above geometric reconstruction target as a dense prediction task. For a given set of N input images, the model is trained to output N corresponding depth maps and ray maps, each pixel-aligned with its respective input. The architecture to achieve this begins with a standard pretrained vision transformer (e.g., Oquab et al. 61), as its backbone, leveraging its powerful feature extraction capabilities. To handle arbitrary view counts, we introduce a key modification: an input-adaptive cross-view self-attention mechanism. This module dynamically rearranges tokens during the forward pass in selected layers, enabling efficient information exchange across all views. For the final prediction, we propose a new dual DPT head designed to jointly outputs both depth and ray values, by processing the same set of features with distinct fusion parameters. To enhance flexibility, the model can optionally incorporate known camera poses via a simple camera encoder, allowing it to adapt to various practical settings. This overall design results in a clean and scalable architecture that directly inherits the scaling properties of its pretrained backbone.

We train Depth Anything 3 via a teacher-student paradigm to unify diverse training data, which is necessary for a generalist model. Our data sources include varied formats like real-world depth camera captures (e.g., Baruch et al. 5), 3D reconstruction (e.g., Reizenstein et al. 68), and synthetic data, where real-world depth may be of poor quality (Fig. 4). To resolve this, we adopt a pseudo-labeling strategy inspired by prior works [112, 113]. Specifically, we train a powerful teacher monocular depth model on synthetic data to generate dense, high-quality pseudo-depth for all real-world data. Crucially, to preserve geometric integrity, we align these dense pseudo-depth maps with the original sparse or noisy depth. This approach proved remarkably effective, significantly enhancing label detail and completeness without sacrificing the geometric accuracy.

To better evaluate our model and track progress in the field, we establish a comprehensive benchmark for assessing geometry and pose accuracy. The benchmark comprises 5 distinct datasets, totaling over 89 scenes, ranging from object-level to indoor and outdoor environments. By directly evaluating pose accuracy across scenes and fusing the predicted pose and depth into a 3D point cloud for accuracy assessment, the benchmark faithfully measures the pose and depth accuracy of visual geometry estimators. Experiments show that our model achieves state-of-the-art performance on 18 out of 20 settings. Moreover, on standard monocular benchmarks, our model outperforms Depth Anything 2 [113].

To further demonstrate the fundamental capability of Depth Anything 3 in advancing other 3D vision tasks, we introduce a challenging benchmark for feed-forward novel view synthesis (FF-NVS), comprising over 160 scenes. We adhere to the minimal modeling strategy and fine-tune our model with an additional DPT head to predict pixel-aligned 3D Gaussian parameters. Extensive experiments yield two key findings: 1) fine-tuning a geometry foundation model for NVS substantially outperforms highly specialized task-specific models [108]; 2) enhanced geometric reconstruction capability directly correlates with improved FF-NVS performance, establishing Depth Anything 3 as the optimal backbone for this task.

2 Related Work

Multi-view visual geometry estimation. Traditional systems [70, 71] decompose reconstruction into feature detection and matching, robust relative pose estimation, incremental or global SfM with bundle adjustment, and dense multi-view stereo for per-view depth and fused point clouds. These methods remain strong on well-textured scenes, but their modularity and brittle correspondences complicate robustness under low texture, specularities, or large viewpoint changes. Early learning methods injected robustness at the component level: learned detectors [20], descriptors for matching [22], and differentiable optimization layers that expose pose/depth updates to gradient flow [31, 33, 62]. On the dense side, cost-volume networks [106, 114] for MVS replaced hand-crafted regularization with 3D CNNs, improving depth accuracy especially at large baselines and thin structures compared with classical PatchMatch. Early end-to-end approaches [86, 90] moved beyond modular SfM/MVS pipelines by directly regressing camera poses and per-image depths from pairs of images.




Figure 2 Pipeline of Depth Anything 3. Depth Anything 3 employs a single transformer (vanilla DINOv2 model) without any architectural modifications. To enable cross-view reasoning, an input-adaptive cross-view self-attention mechanism is introduced. A dual-DPT head is used to predict depth and ray maps from visual tokens. Camera parameters, if available, are encoded as camera tokens and concatenated with patch tokens, participating in all attention operations.

These approaches reduced engineering complexity and demonstrated the feasibility of learned joint depth pose estimation, but they often struggled with scalability, generalization, and handling arbitrary input cardinalities.

A turning point came with DUST3R [96], which leveraged transformers to directly predict point map between two views and compute both depth and relative pose in a purely feed-forward manner. This work laid the foundation for subsequent transformer-based methods aiming to unify multi-view geometry estimation at scale. Follow-up models extended this paradigm with multi-view inputs [10, 85, 94, 110], video input [19, 59, 94, 121], robust correspondence modeling [48], camera parameter injection [39, 43], large-scale SfM [18], SLAM applications [54], and view synthesis with 3D Gaussians [11, 13, 41, 79, 108, 122]. Among these, [91] push accuracy to a new level through large-scale training, a multi-stage architecture, and redundancy in design. In contrast, we focus on a minimal modeling strategy built around a single, simple transformer.

Monocular depth estimation. Early monocular depth estimation methods relied on fully supervised learning on single-domain datasets, which often produced models specialized to either indoor rooms [75] or outdoor driving scenes [26]. These early deep models achieved good accuracy within their training domain but struggled to generalize to novel environments, highlighting the challenge of cross-domain depth prediction. Modern generalist approaches [6, 42, 95, 112, 113, 118] exemplify this trend by leveraging massive multi-dataset training and advanced architectures like vision transformers [67] or DiT [64]. Trained on millions of images, they learn broad visual cues and incorporate techniques such as affine-invariant depth normalization. In contrast, our method is primarily designed for a unified visual geometry estimation task, yet it still demonstrates competitive monocular depth performance.

Feed-Forward Novel View Synthesis Novel view synthesis (NVS) has long been a core problem in computer vision and graphics [8, 34, 49], and interest has increased with the rise of neural rendering [28, 44, 57, 77, 78]. A particularly promising direction is feed-forward NVS, which produces 3D representations in a single pass through an image-to-3D network, avoiding tedious per-scene optimization. Early methods adopted NeRF as the underlying 3D representation [12, 14, 35, 52, 107, 119], but recent work has largely shifted to 3DGS due to its explicit structure and real-time rendering. Representative approaches improve image-to-3D networks with geometry priors, *e.g.*, epipolar attention [11], cost volumes [13], and depth priors [107]. More recently, multi-view geometry foundation models [85, 91, 96, 110] have been integrated to improve modeling capacity, particularly in pose-free settings, yet methods built upon such models are often evaluated by relying on a single chosen foundation model [41, 79, 116]. Here, we systematically benchmark the contribution of different geometry foundation models to NVS and propose strategies to better exploit them, enabling feed-forward 3DGS to handle both posed and pose-free inputs, variable numbers of views, and arbitrary resolutions.

3 Depth Anything 3

We tackle the recovery of consistent 3D geometry from diverse visual inputs—single image, multi-view collections, or videos—and optionally incorporate known camera poses when available.

3.1 Formulation

We denote the input as $\mathcal{I} = \{\mathbf{I}_i\}_{i=1}^{N_v}$ with each $\mathbf{I}_i \in \mathbb{R}^{H \times W \times 3}$. For $N_v = 1$ this is a monocular image, and for $N_v > 1$ it represents a video or multi-view set. Each image has depth $\mathbf{D}_i \in \mathbb{R}^{H \times W}$, camera extrinsics $[\mathbf{R}_i \mid \mathbf{t}_i]$, and intrinsics $\mathbf{K}_i$. The camera can also be represented as $\mathbf{v}_i \in \mathbb{R}^9$ with translation $\mathbf{t}_i \in \mathbb{R}^3$, rotation quaternion $\mathbf{q}_i \in \mathbb{R}^4$, and FOV parameters $\mathbf{f}_i \in \mathbb{R}^2$. A pixel $\mathbf{p} = (u, v, 1)^\top$ projects to a 3D point $\mathbf{P} = (X, Y, Z, 1)^\top$ by

$$\mathbf{P} = \mathbf{R}_i(\mathbf{D}_i(u, v) \mathbf{K}_i^{-1} \mathbf{p}) + \mathbf{t}_i,$$

through which the underlying 3D visual space can be faithfully recovered.

Depth-ray representation. Predicting a valid rotation matrix $\mathbf{R}_i$ is challenging due to the orthogonality constraint. To avoid this, we represent camera pose implicitly with a per-pixel ray map, aligned with the input image and depth map. For each pixel $\mathbf{p}$, the camera ray $\mathbf{r} \in \mathbb{R}^6$ is defined by its origin $\mathbf{t} \in \mathbb{R}^3$ and direction $\mathbf{d} \in \mathbb{R}^3$: $\mathbf{r} = (\mathbf{t}, \mathbf{d})$. The direction is obtained by backprojecting $\mathbf{p}$ into the camera frame and rotating it to the world frame: $\mathbf{d} = \mathbf{R} \mathbf{K}^{-1} \mathbf{p}$. The dense ray map $\mathbf{M} \in \mathbb{R}^{H \times W \times 6}$ stores these parameters for all pixels. We do not normalize $\mathbf{d}$, so its magnitude preserves the projection scale. Thus, a 3D point in world coordinates is simply $\mathbf{P} = \mathbf{t} + \mathbf{D}(u, v) \cdot \mathbf{d}$. This formulation enables consistent point cloud generation by combining predicted depth and ray maps through element-wise operations.

Deriving Camera Parameters from the Ray Map. Given an input image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, the corresponding ray map is denoted by $\mathbf{M} \in \mathbb{R}^{H \times W \times 6}$. This map comprises per-pixel ray origins, stored in the first three channels ($\mathbf{M}(:, :, : 3)$), and ray directions, stored in the last three ($\mathbf{M}(:, :, 3 :)$).

First, the camera center $\mathbf{t}_c$ is estimated by averaging the per-pixel ray origin vectors:

$$\mathbf{t}_c = \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W \mathbf{M}(h, w, : 3). \quad (1)$$

To estimate the rotation $\mathbf{R}$ and intrinsics $\mathbf{K}$, we formulate the problem as finding a homography $\mathbf{H}$. We begin by defining a “identity” camera with an identity intrinsics matrix, $\mathbf{K}_I = \mathbf{I}$. For a given pixel $\mathbf{p}$, its corresponding ray direction in this canonical camera’s coordinate system is simply $\mathbf{d}_I = \mathbf{K}_I^{-1} \mathbf{p} = \mathbf{p}$. The transformation from this canonical ray $\mathbf{d}_I$ to the ray direction $\mathbf{d}_{\text{cam}}$ in the target camera’s coordinate system is given by $\mathbf{d}_{\text{cam}} = \mathbf{K} \mathbf{R} \mathbf{d}_I$. This establishes a direct homography relationship, $\mathbf{H} = \mathbf{K} \mathbf{R}$, between the two sets of rays. We can then solve for this homography by minimizing the geometric error between the transformed canonical rays and a set of pre-computed target rays, $\mathbf{M}(h, w, 3 :)$. This leads to the following optimization problem:

$$\mathbf{H}^* = \arg \min_{\|\mathbf{H}\|=1} \sum_{h=1}^H \sum_{w=1}^W \|\mathbf{H} \mathbf{p}_{h,w} \times \mathbf{M}(h, w, 3 :)\|. \quad (2)$$

This is a standard least-squares problem that can be efficiently solved using the Direct Linear Transform (DLT) algorithm [2]. Once the optimal homography $\mathbf{H}^*$ is found, we recover the camera parameters. Since the intrinsic matrix $\mathbf{K}$ is upper-triangular and the rotation matrix $\mathbf{R}$ is orthonormal, we can uniquely decompose $\mathbf{H}^*$ using RQ decomposition to obtain $\mathbf{K}$, $\mathbf{R}$.

Minimal prediction targets. Recent works aim to build unified models for diverse 3D tasks, often using multitask learning with different targets—for example, point maps alone [96], or redundant combinations of pose, local/global point maps, and depth [91, 94, 110]. While point maps are insufficient to ensure consistency, redundant targets can improve pose accuracy but often introduce entanglement that compromises it. In

contrast, our experiments (Tab. 6) show that a depth-ray representation forms a minimal yet sufficient target set for capturing both scene structure and camera motion, outperforming alternatives like point maps or more complex outputs. However, recovering camera pose from the ray map at inference is computationally costly. We address this by adding a lightweight camera head, $\mathcal{D}_C$. This transformer operates on camera tokens to predict the field of view ($\mathbf{f} \in \mathbb{R}^2$), rotation as a quaternion ($\mathbf{q} \in \mathbb{R}^4$), and translation ($\mathbf{t} \in \mathbb{R}^3$). Since it processes only one token per view, the added cost is negligible.

3.2 Architecture

We now detail the architecture of Depth Anything 3, which is illustrated in Fig. 2. The network is composed of three main components: a single transformer model as the backbone, an optional camera encoder for pose conditioning, and a Dual-DPT head for generating predictions.

Single transformer backbone. We use a Vision Transformer with L blocks, pretrained on large-scale monocular image corpora (e.g., DINOv2 [61]). Cross-view reasoning is enabled without architectural changes via an input-adaptive self-attention, implemented by rearranging input tokens. We divide the transformer into two groups of sizes L_s and L_g . The first L_s layers apply self-attention within each image, while the subsequent L_g layers alternate between cross-view and within-view attention, operating on all tokens jointly through tensor reordering. In practice, we set $L_s : L_g = 2 : 1$ with $L = L_s + L_g$. As shown in our ablation study in Tab. 7, this configuration provides the optimal trade-off between performance and efficiency compared to other arrangements. This design is input-adaptive: with a single image, the model naturally reduces to monocular depth estimation without extra cost.

Camera condition injection. To seamlessly handle both posed and unposed inputs, we prepend each view with a camera token $\mathbf{c}_i$. If camera parameters ($\mathbf{K}_i, \mathbf{R}_i, \mathbf{t}_i$) are available, the token is obtained via a lightweight MLP $\mathcal{E}_c$: $\mathbf{c}_i = \mathcal{E}_c(\mathbf{f}_i, \mathbf{q}_i, \mathbf{t}_i)$. Otherwise, a shared learnable token $\mathbf{c}_l$ is used. Concatenated with patch tokens, these camera tokens participate in all attention operations, providing either explicit geometric context or a consistent learned placeholder.

Dual-DPT head. For the final prediction stage, we propose a novel Dual-DPT head that jointly produces dense depth and ray values. As shown in Tab. 6, this design is both powerful and efficient. Given a set of features from the backbone, the Dual-DPT head first processes them through a shared set of reassembly modules. Subsequently, the processed features are fused using two distinct sets of fusion layers: one for the depth branch and one for the ray branch. Finally, two separate output layers produce the final depth and ray map predictions. This architecture ensures that both branches operate on the same set of processed features, differing only in the final fusion stage. Such a design encourages strong interaction between the two prediction tasks, while avoiding redundant intermediate representations.




Figure 3 Dual-DPT Head. Two branches share reassembly modules for better outputs alignment.

3.3 Training

Teacher-student learning paradigm. Our training data comes from diverse sources, including real-world depth captures, 3D reconstructions, and synthetic datasets. Real-world depth is often noisy and incomplete (Fig. 4), limiting its supervisory value. To mitigate this, we train a monocular relative depth estimation “teacher” model solely on synthetic data to generate high-quality pseudo-labels. These pseudo-depth maps are aligned with the original sparse or noisy ground truth via RANSAC least squares, enhancing label detail and completeness while preserving geometric accuracy. We term this model Depth-Anything-3-Teacher, trained on

a large synthetic corpus covering indoor, outdoor, object-centric, and diverse in-the-wild scenes to capture fine geometry. We detail our teacher design in the Sec. 4.1.

Training objectives. Following the formulation in Sec. 3.1, our model $\mathcal{F}_\theta$ maps an input $\mathcal{I}$ to a set of outputs comprising a depth map $\hat{\mathbf{D}}$, a ray map $\hat{\mathbf{R}}$, and an optional camera pose $\hat{\mathbf{c}}$: $\mathcal{F}_\theta : \mathcal{I} \mapsto \{\hat{\mathbf{D}}, \hat{\mathbf{R}}, \hat{\mathbf{c}}\}$. The gray color indicates that $\hat{\mathbf{c}}$ is an optional output, included primarily for practical convenience. Prior to loss computation, all ground-truth signals are normalized by a common scale factor. This scale is defined as the mean ℓ_2 norm of the valid reprojected point maps $\mathbf{P}$, a step that ensures consistent magnitude across different modalities and stabilizes the training process. The overall training objective is defined as a weighted sum of several terms:

$$\mathcal{L} = \mathcal{L}_D(\hat{\mathbf{D}}, \mathbf{D}) + \mathcal{L}_M(\hat{\mathbf{R}}, \mathbf{M}) + \mathcal{L}_P(\hat{\mathbf{D}} \odot \mathbf{d} + \mathbf{t}, \mathbf{P}) + \beta \mathcal{L}_C(\hat{\mathbf{c}}, \mathbf{v}) + \alpha \mathcal{L}_{\text{grad}}(\hat{\mathbf{D}}, \mathbf{D}),$$

$$\mathcal{L}_D(\hat{\mathbf{D}}, \mathbf{D}; D_c) = \frac{1}{Z_\Omega} \sum_{p \in \Omega} m_p \left(D_{c,p} |\hat{\mathbf{D}}_p - \mathbf{D}_p| - \lambda_c \log D_{c,p} \right),$$

where $D_{c,p}$ denotes the confidence of depth D_p . All loss terms are based on the ℓ_1 norm, with weights set to $\alpha = 1$ and $\beta = 1$. The gradient loss, $\mathcal{L}_{\text{grad}}$, penalizes the depth gradients:

$$\mathcal{L}_{\text{grad}}(\hat{\mathbf{D}}, \mathbf{D}) = \|\nabla_x \hat{\mathbf{D}} - \nabla_x \mathbf{D}\|_1 + \|\nabla_y \hat{\mathbf{D}} - \nabla_y \mathbf{D}\|_1, \quad (3)$$

where ∇_x and ∇_y are the horizontal and vertical finite difference operators. This loss preserves sharp edges while ensuring smoothness in planar regions. In practice, we set $\alpha = 1$ and $\beta = 1$.

3.4 Implementation Details

Training datasets. We provide our training datasets in Table 1. Note that for datasets with potential overlap between training and testing (ScanNet++), we ensure a strict separation at the scene level, i.e., scenes in training and testing are mutually exclusive. Note that using Scannet++ for training is fair to other methods, as it is widely used for training in [91, 96].

Training details. We train our model on 128 H100 GPUs for 200k steps, using an 8k-step warm-up and a peak learning rate of 2×10^{-4} . The base resolution is 504×504 , which is divisible by 2, 3, 4, 6, 9 and 14, making it more compatible with common photo aspect ratios such as 2:3, 3:4, and 9:16. Training image resolutions are randomly sampled from 504×504 , 504×378 , 504×336 , 504×280 , 336×504 , 896×504 , 756×504 , 672×504 . For the 504×504 resolution, the number of views is sampled uniformly from [2, 18]. The batch size is dynamically adjusted to keep the token count per step approximately constant. Supervision transitions from ground-truth depth to teacher-model labels at 120k steps. Pose conditioning is randomly activated during training with probability 0.2.

4 Teacher-Student Learning

As shown in Fig. 4, the real-world datasets are of poor quality, thus we train the teacher model exclusively on synthetic data to provide supervision for real-world data. Our teacher model is trained as a monocular relative depth predictor. During inference or supervision, noisy ground-truth depth can be used to provide scale and shift parameters, allowing for the alignment of the predicted relative depth with absolute depth measurements.

4.1 Constructing the Teacher Model

Building upon Depth Anything 2 [112], we extend the approach in several key aspects, including both




Figure 4 Poor quality real-world datasets. We show some examples of the poor quality real-world datasets.

Table 1 Datasets used in Depth Anything 3 , including number of scenes, data type.

Usage	Dataset	#Scenes	Data Type
Pose-geometry benchmark	HiRoom (ours)	29	Synthetic
	ETH3D [72]	11	LiDAR
	DTU [1]	22	LiDAR
	7Scenes [74]	7	LiDAR
	ScanNet++ [117]	20	LiDAR
Pose-geometry Training	AriaDigitalTwin [63]	237	Synthetic
	AriaSyntheticENV [63]	99950	Synthetic
	ArkitScenes [5]	4388	LiDAR
	BlendedMVS [115]	503	3D Recon
	Co3dv2 [68]	30616	Colmap
	DL3DV [53]	6379	Colmap
	HyperSim [69]	344	Synthetic
	MapFree [3]	921	Colmap
	MegaDepth [51]	268	Colmap
	MegaSynth [40]	6049	Synthetic
	MvsSynth [38]	121	Synthetic
	Objaverse [17]	505557	Synthetic
	Omniobject [102]	5885	Synthetic
	OmniWorld [128]	1039	Synthetic
	PointOdyssey [127]	44	Synthetic
	ReplicaVMAP [83]	17	Synthetic
	ScanNet++ [117]	230	LiDAR
	ScenenetRGBD [55]	16866	Synthetic
	TartanAir [98]	355	Synthetic
	Trellis [104]	557408	Synthetic
	vKitti2 [9]	50	Synthetic
	WildRGBD [103]	23050	LiDAR

data and representation. We observe that expanding the training corpus yields clear improvements in depth estimation performance, substantiating the benefits of data scaling. Furthermore, while our revised depth representation may not show striking improvements on standard 2D evaluation metrics, it leads to qualitatively better 3D point clouds, exhibiting fewer geometric distortions and more realistic scene structures. It is important to note that our teacher network backbone is directly aligned with the above DA3 framework, consisting solely of a DINOv2 vision transformer with a DPT decoder—no specialized architectural modifications are introduced. We elaborate on the full design and implementation details in the following sections.

Data scaling. We train the teacher model exclusively on synthetic data to achieve finer geometric detail. The synthetic datasets used in DA2 are relatively limited. In DA3, we substantially expand the training corpus to include: Hypersim [69], TartanAir [98], IRS [93], vKITTI2 [9], BlendedMVS [115], SPRING [56], MVSSynth [38], UnrealStereo4K [123], GTA-SfM [92], TauAgent [27], KenBurns [60], MatrixCity [50], EDEN [47], ReplicaGSO [82], UrbanSyn [32], PointOdyssey [127], Structured3D [125], Objaverse [17], Trellis [104], and OmniObject [102]. This collection spans indoor, outdoor, object-centric, and diverse in-the-wild scenes, improving generalization of the teacher model.

Depth representation. Unlike DA2, which predicts scale-shift-invariant disparity, our teacher outputs scale-shift-invariant depth. Depth is preferable for downstream tasks, such as metric depth estimation and multiview geometry, that directly operate in depth space rather than disparity. To address depth’s reduced sensitivity for near-camera regions comparing to disparity, we predict exponential depth instead of linear depth, enhancing discrimination at small distances.

Training objectives. For geometric supervision, in addition to a standard depth-gradient loss, we adopt ROE alignment with the global-local loss introduced in [95]. To further refine local geometry, we introduce a distance-weighted surface-normal loss. For each center pixel, we sample four neighboring points and compute unnormalized normals n_i . We then weight these normals by:

$$w_i = \sum_{j=0}^4 \|n_j\| - \|n_i\|, \quad (4)$$

which downweights contributions from neighbors farther from the center, yielding a mean normal closer to the true local surface normal:

$$n_m = \sum_{i=0}^4 w_i \frac{n_i}{\|n_i\|}, \quad (5)$$

The final normal loss is

$$\mathcal{L}_N = \mathcal{E}(\hat{n}_m, n_m) + \sum_{i=0}^4 \mathcal{E}(\hat{n}_i, n_i) \quad (6)$$

where $\mathcal{E}$ denotes the angular error between normals. Ground truth is undefined in sky regions and in background areas of object-only datasets. To prevent these regions from degrading the depth prediction and to facilitate downstream use, we jointly predict a sky mask and an object mask aligned with the depth output, supervised with MSE loss. The overall training objective is

$$\mathcal{L}_T = \alpha \mathcal{L}_{\text{grad}} + \mathcal{L}_{\text{gl}} + \mathcal{L}_N + \mathcal{L}_{\text{sky}} + \mathcal{L}_{\text{obj}} \quad (7)$$

where $\alpha = 0.5$. Here, $\mathcal{L}_{\text{grad}}$, $\mathcal{L}_{\text{gl}}$, $\mathcal{L}_{\text{sky}}$, and $\mathcal{L}_{\text{obj}}$ denote the gradient loss, global-local loss, sky-mask loss, and object-mask loss, respectively.

4.2 Teaching Depth Anything 3

Real-world datasets are crucial for generalizing camera pose estimation, yet they rarely provide clean depths; supervision is often noisy or sparse (Fig. 4). Depth Anything 3 Teacher provides high-quality relative depth, which we align to noisy metric measurements (e.g., COLMAP or active sensors) via a robust ransac scale-shift procedure. Let $\tilde{\mathbf{D}}$ denote the teacher’s relative depth and $\mathbf{D}$ the available sparse depth with validity mask m_p over domain Ω . We estimate scale s and shift t by RANSAC least squares, using an inlier threshold equal to the mean absolute deviation from the residual median:

$$(\hat{s}, \hat{t}) = \arg \min_{s>0, t} \sum_{p \in \Omega} m_p (s \tilde{\mathbf{D}}_p + t - \mathbf{D}_p)^2, \quad \mathbf{D}^{T \rightarrow M} = \hat{s} \tilde{\mathbf{D}} + \hat{t}. \quad (8)$$

The aligned $\mathbf{D}^{T \rightarrow M}$ provides scale-consistent and pose-depth coherent supervision for Depth Anything 3, complementing our joint depth-ray objectives and improving real-world generalization, as evidenced in Fig. 8.

4.3 Teaching Monocular Model

We additionally train a monocular depth model under a teacher-student paradigm. We follow the DA2 framework, training the monocular student on unlabeled images with teacher-generated pseudo-labels. The key difference from DA2 lies in the prediction target: our student predicts depth maps, whereas DA2 predicts disparity. We further supervise the student with the same loss used for the teacher, applied to the pseudo-depth labels. The monocular model also predicts relative depth. Trained solely on unlabeled data with teacher supervision, it achieves state-of-the-art performance on standard monocular depth benchmarks as shown in Tab. 10.

4.4 Teaching Metric Model

Next, we demonstrate that our teacher model can be used for training a metric depth estimation model with sharp boundaries. Following Metric3Dv2 [37], we apply canonical camera space transformation to address

depth ambiguity caused by varying focal lengths. Specifically, we rescale the ground-truth depth using ratio f^c/f , where f^c and f denote the canonical focal length and camera focal length, respectively. To ensure sharp details, we employ Teacher model’s prediction as training labels. We align the scale and shift of the teacher model’s predicted depths with the ground-truth metric depth labels for supervision.

Training dataset. We trained our metric depth model on 14 datasets, including Taskonomy [120], DIML (Outdoor) [15], DDAD [30], Argoverse [101], Lyft [], PandaSet [105], Waymo [84], ScanNet++ [117], ARKitScenes [5], Map-free [3], DSEC [24], Driving Stereo [109] and Cityscapes [16] datasets. For stereo datasets, we leverage the prediction of FoundationStereo [100] as training labels.

Implementation Details. The training largely follows that of the monocular teacher model. All images are trained at a base resolution of 504 with varying aspect ratios (1:1, 1:2, 16:9, 9:16, 3:4, 1:1.5, 1.5:1, 1:1.8). We employ AdamW optimizer and set the learning rate for encoder and decoder to $5e-6$ and $5e-5$, respectively. We apply random rotation augmentation where training images are rotated at 90 or 270 degree with 5% probability. We set canonical focal length f^c to 300. We use the aligned prediction from teacher model as supervision. With a probability of 20%, we use the original ground-truth labels for training. We train with batch size of 64 for 160K iterations. The training objective is a weighted sum of depth loss $\mathcal{L}_{depth}$, $\mathcal{L}_{grad}$ and sky-mask loss $\mathcal{L}_{sky}$.

5 Application: Feed-Forward 3D Gaussian Splattings

5.1 Pose-Conditioned Feed-Forward 3DGS

Inspired by human spatial intelligence, we believe that consistent depth estimation can greatly enhance downstream 3D vision tasks. We choose feed-forward novel view synthesis (FF-NVS) as the demonstration task, given its growing attention driven by advances in neural 3D representations (*i.e.*, we choose 3DGS) and its relevance to numerous applications. Adhere to the minimal modeling strategy, we perform FF-NVS by *fine-tuning* with an added DPT head (GS-DPT) to infer pixel-aligned 3D Gaussians [11, 13].

GS-DPT head. Given visual tokens for each view extracted via our single transformer backbone (Sec. 3.2), GS-DPT predicts the camera-space 3D Gaussian parameters $\{\sigma_i, \mathbf{q}_i, \mathbf{s}_i, \mathbf{c}_i\}_{i=1}^{H \times W}$, where $\sigma_i, \mathbf{q}_i \in \mathbb{H}$, $\mathbf{s}_i \in \mathbb{R}^3$, $\mathbf{c}_i \in \mathbb{R}^3$ denote the opacity, rotation quaternion, scale, and RGB color of the i -th 3D Gaussian, respectively. Among them, σ_i is predicted by the confidence head, while others are predicted by the main GS-DPT head. The estimated depth is unprojected to world coordinates to obtain the global positions $\mathbf{P}_i \in \mathbb{R}^3$ of the 3D Gaussians. These primitives are then rasterized to synthesize novel views from given camera poses.

Training objectives. The NVS model is fine-tuned with two training objectives, namely photometric loss (*i.e.*, $\mathcal{L}_{MSE}$ and $\mathcal{L}_{LPIPS}$) on rendered novel views and scale-shift-invariant depth loss $\mathcal{L}_D$ on the estimated depth of observed views, following the teacher-student learning paradigm (Sec. 3.3).

5.2 Pose-Adaptive Feed-Forward 3DGS

Unlike the above pose-conditioned version intended to benchmark DA3 as a strong feed-forward 3DGS backbone, we also present an alternative better suited to in-the-wild evaluation. This version is designed to integrate seamlessly with DA3 using *identical* pretrained weights, enabling novel view synthesis with or without camera poses, and across varying resolutions and input view counts.

Pose-adaptive formulation. Rather than assuming that all input images are uncalibrated [41, 79, 116, 122], we adopt a *pose-adaptive* design that accepts both posed and unposed inputs, yielding a flexible framework that works *with or without poses*. Two design choices are required to achieve this: 1) all 3DGS parameters are predicted in local camera space. 2) the backbone must handle posed and unposed images seamlessly. Our DA3 backbone satisfies both requirements (Sec. 3.2). In particular, when poses are available, we scale (via [87]) and unproject the predicted depth and camera-space 3DGS to world space to align with them. When poses are not available, we directly use the predicted poses for the unprojection to world space.

To reduce the trade-off between accurate surface geometry and rendering quality [29], we predict an additional depth offset in the GS-DPT head. For more in-the-wild robustness, we replace per 3D Gaussian color with spherical harmonic coefficients to reduce conflicts with geometry via modeling view-dependent surface.

Enhanced training strategies. To avoid unstable training, we initialize the DA3 backbone from pretrained weights and *freeze* it when training, tuning only the GS-DPT head. To improve in-the-wild performance, we train with varying image resolutions and varying numbers of context views. Specifically, higher-resolution inputs are paired with fewer context views and lower-resolution inputs with more views, which stabilizes training while supporting diverse evaluation scenarios.

5.3 Implementation Details

Training datasets. For training the NVS model, we leverage the large-scale DL3DV dataset [53], which provides diverse real-world scenes with camera poses estimated by COLMAP. We use 10,015 scenes from DL3DV for training the feed-forward 3DGS model. To ensure fair evaluation, we strictly maintain exclusivity between training and testing splits: the 140 DL3DV scenes used for benchmarking are completely disjoint from the training set, preventing any data leakage.

6 Visual Geometry Benchmark

We further introduce a visual geometry benchmark to assess geometry prediction models. It directly evaluates pose accuracy, depth via reconstruction accuracy and visual rendering quality.

6.1 Benchmark Pipeline

Pose estimation. For each scene, we select all available images; if the total number exceeds the limit, we randomly sample 100 images using a fixed random seed. The selected images are then processed through a feed-forward model to generate consistent pose and depth estimations, after which the pose accuracy is computed.

Geometry estimation. For the same image set, we perform reconstruction using the predicted poses together with the predicted depths. To align the reconstructed point cloud with the ground-truth, we employ *evo* [87] to align the predicted poses to the ground-truth poses, obtaining a transformation that maps the reconstruction into the ground-truth coordinate system. To improve robustness, we adopt a RANSAC-based alignment procedure. Specifically, we repeatedly apply *evo* on randomly sampled pose subsets and evaluate each candidate transformation by counting the number of inlier poses, where inliers are defined as those with translation errors below the median of the overall pose deviations. The transformation with the largest inlier set is then chosen and applied to fuse the aligned predicted point cloud with the predicted depth maps by TSDF fusion. Finally, reconstruction quality is assessed by comparing the aligned reconstruction with the ground-truth point cloud using the metrics described in Sec. 6.2.

Visual rendering. For each testing scene, the number of images typically ranges from 300 to 400 across all benchmark datasets. We sample one out of every 8 images as target novel views for evaluation. From the remaining viewpoints, we use COLMAP camera poses provided by each dataset and apply farthest point sampling, considering both camera translation and rotation distances, to select 12 images as input context views. For DL3DV, we use the official Benchmark set for testing. For Tanks and Temples, all Training Data scenes are included except *Courthouse*. For MegaDepth, we select scenes numbered from 5000 to 5018, as these are most suitable for NVS.

6.2 Metrics

Pose metrics. For assessing pose estimation, we follow the evaluation protocol introduced in [89, 91] and report results using the AUC. This metric is derived from two components: Relative Rotation Accuracy (RRA) and Relative Translation Accuracy (RTA). RRA and RTA quantify the angular deviation in rotation and translation, respectively, between two images. Each error is compared against a set of thresholds to obtain

accuracy values. AUC is then computed as the integral of the accuracy–threshold curve, where the curve is determined by the smaller of RRA and RTA at each threshold. To illustrate performance under different tolerance levels, we primarily report results at thresholds of 3 and 30.

Reconstruction metrics. Let $\mathcal{G}$ denote the ground-truth point set and $\mathcal{R}$ the reconstructed point set under evaluation. We measure accuracy using $\text{dist}(\mathcal{R} \rightarrow \mathcal{G})$ and completeness using $\text{dist}(\mathcal{G} \rightarrow \mathcal{R})$ following [1]. The Chamfer Distance (CD) is then defined as the average of these two terms. Based on these distances, we define the precision and recall of the reconstruction $\mathcal{R}$ with respect to a distance threshold d . Precision is given by $\frac{1}{|\mathcal{R}|} \sum [\text{dist}(\mathcal{R}_i \rightarrow \mathcal{G}) < d]$, and recall by $\frac{1}{|\mathcal{G}|} \sum [\text{dist}(\mathcal{G}_i \rightarrow \mathcal{R}) < d]$, where $[\cdot]$ denotes the Iverson bracket [46]. To jointly capture both measures, we report the F1-score, computed as $\text{F1} = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$.

6.3 Datasets

Our benchmark is built on five datasets: HiRoom [129], ETH3D [72], DTU [1], 7Scenes [74], and ScanNet++ [117]. Together, they cover diverse scenarios ranging from object-centric captures to complex indoor and outdoor environments, and are widely adopted in prior research. Below, we present more details about the dataset preparation process.

HiRoom is a Blender-rendered synthetic dataset comprising 30 indoor living scenes created by professional artists. We use a threshold d of 0.05m for the F1 reconstruction metric calculation. For TSDF fusion, we set the parameters voxel size to 0.007m.

ETH3D provides high-resolution indoor and outdoor images with ground-truth depth from laser sensors. We aggregate the ground-truth depth maps with TSDF fusion for GT 3D shapes. We select 11 scenes: *courtyard*, *electro*, *kicker*, *pipes*, *relief*, *delivery area*, *facade*, *office*, *playground*, *relief 2*, *terrains*, for the benchmark. All frames are used in the evaluation. We use a threshold d of 0.25 for the F1 reconstruction metric calculation. For TSDF fusion, we set the parameters voxel size to 0.039m.

DTU is an indoor dataset consisting of 124 different objects, each scene is recorded from 49 views. It provides ground-truth point clouds collected under well-controlled conditions. We evaluate models on the 22 evaluation scans of the DTU dataset following [114]. We adopt the RMBG 2.0 [126] to remove meaningless background pixels and use the default depth fusion strategy proposed in [124]. All frames are used in the evaluation.

7Scenes is a challenging real-world dataset, consisting of low-resolution images with severe motion blurs for in-door scenes. We follow the implementation in [130] to fuse RGBD images with TSDF fusion and prepare ground-truth 3D shapes. We downsample the number of frames for each scene by 11 to facilitate evaluation. We use a threshold d of 0.05m for the F1 reconstruction metric calculation. For TSDF fusion, we set the parameters voxel size to 0.007m.

ScanNet++ is an extensive indoor dataset providing high-resolution images, depth maps from iPhone LiDAR, and high-resolution depth maps sampled from reconstructions of laser scans. We select 20 scenes for the benchmark. As depth maps from iPhone LiDAR lack of invalid ground-truth indicators, we use depth maps sampled from reconstructions of laser scans as ground-truth depth by default. We aggregate the ground-truth depth maps with TSDF fusion for GT 3D shapes. We downsample the number of frames for each scene by 5 to facilitate evaluation. We use a threshold d of 0.05m for the F1 reconstruction metric calculation. For TSDF fusion, we set the parameters voxel size to 0.02m.

Visual rendering quality. We evaluate visual rendering quality on diverse large-scale scenes. We introduce a new NVS benchmark built from three datasets, including DL3DV [53] with 140 scenes, Tanks and Temples [45] with 6, and MegaDepth [51] with 19, each spanning around 300 sampled frames. Ground truth camera poses, estimated with COLMAP, are used directly to ensure accurate and fair comparison across diverse models. We report PSNR, SSIM, and LPIPS metrics on rendered novel views using given camera poses.

Table 2 Comparisons with SOTA methods on pose accuracy. We report both Auc3 $\uparrow$ and Auc30 $\uparrow$ metrics. The top-3 results are highlighted as first, second, and third.

Methods	Params	HiRoom		ETH3D		DTU		7Scenes		ScanNet++	
		Auc3	Auc30	Auc3	Auc30	Auc3	Auc30	Auc3	Auc30	Auc3	Auc30
DUSt3R	0.57B	17.6	54.3	4.30	27.3	4.00	74.3	6.90	61.6	8.10	33.9
Fast3R	0.65B	25.9	77.0	8.10	44.4	9.50	79.1	19.0	78.6	17.9	72.5
MapAnything	0.56B	17.9	82.8	19.2	77.4	6.50	72.7	12.6	79.7	20.2	84.1
Pi3	0.96B	67.0	94.8	35.2	87.3	62.5	94.9	25.5	86.3	50.7	92.1
VGGT	1.19B	49.1	88.0	26.3	80.8	79.2	99.8	23.9	85.0	62.6	95.1
DA3-Giant	1.10B	80.3	95.9	48.4	91.2	94.1	99.4	28.5	86.8	85.0	98.1
DA3-Large	0.36B	58.7	94.2	32.2	86.9	70.2	96.7	29.2	86.6	60.2	94.7
DA3-Base	0.11B	19.0	83.2	15.1	74.6	60.1	95.9	20.1	82.9	25.1	83.4
DA3-Small	0.03B	9.49	75.2	8.59	62.1	30.6	91.2	14.0	78.7	10.9	71.9




Figure 5 Comparisons of pose estimation quality. Camera trajectories for two videos are shown. Ground-truth trajectories are derived using COLMAP on images with dynamic objects masked.

7 Experiments

7.1 Comparison with State of the Art

Baselines. VGGT [91] is an end-to-end transformer that jointly predicts camera parameters, depth, and 3D points from one or many views. Pi3 [99] further adopts a permutation-equivariant design to recover affine-invariant cameras and scale-invariant point maps from unordered images. MapAnything [43] provides a feed-forward framework that can also take camera pose as input for dense geometric prediction. Fast3R [111] extends point-map regression to hundreds or even thousands of images in a single forward pass. Finally, DUSt3R [97] tackles uncalibrated image pairs by regressing point maps and aligning them globally. Our method is similar to VGGT [91], but adopts a new architecture and a different camera representation, and it is orthogonal to Pi3 [99].

Pose estimation. As shown in Tab. 2 and Fig. 5, comparing against five baselines [43, 91, 96, 99, 110], our DA3-Giant model attains the best performance on nearly all metrics, with the only exception being Auc30 on the DTU dataset. Notably, on Auc3 our model delivers at least an 8% relative improvement over all competing methods, and on ScanNet++ it achieves a 33% relative gain over the second-best model.

Geometry estimation. As shown in Tab. 3, our DA3-Giant establishes a new SOTA in nearly all scenarios, outperforming all competitors in all five pose-free settings. On average, DA3-Giant achieves a relative

Table 3 Comparisons with SOTA methods on reconstruction accuracy. For all datasets except DTU, we report the F-Score (**F1** $\uparrow$). For DTU, we report the chamfer distance (**CD** $\downarrow$, unit: mm). w/o p. and w/ p. denote without pose and with pose, indicating whether ground-truth camera poses are provided for reconstruction. The top-3 results are highlighted as **first**, **second**, and **third**.

Methods	Params	HiRoom		ETH3D		DTU		7Scenes		ScanNet++	
		w/o p.	w/ p.	w/o p.	w/ p.	w/o p.	w/ p.	w/o p.	w/ p.	w/o p.	w/ p.
DUSt3R	0.57B	30.1	39.5	19.7	18.8	7.60	7.97	26.6	39.8	18.9	27.3
Fast3R	0.65B	40.7	48.2	38.5	50.3	6.88	8.20	41.0	49.8	37.1	53.7
MapAnything	0.56B	32.4	69.2	54.8	71.9	7.91	3.97	44.8	55.2	39.4	71.3
Pi3	0.96B	75.8	85.0	72.7	80.6	3.28	1.72	44.2	57.5	63.1	73.3
VGGT	1.19B	56.7	70.2	57.2	66.7	2.05	1.44	47.9	51.4	66.4	70.7
DA3-Giant	1.10B	85.1	95.6	79.0	87.1	1.85	1.85	53.5	56.5	77.0	79.3
DA3-Large	0.36B	69.5	87.1	65.8	75.2	2.08	1.23	56.3	49.2	67.9	75.7
DA3-Base	0.11B	25.9	71.4	49.5	66.7	2.87	2.36	49.9	50.6	47.2	67.8
DA3-Small	0.03B	18.3	52.2	41.6	63.4	5.83	2.49	41.0	46.8	32.3	53.8

improvement of **25.1%** over VGGT and **21.5%** over Pi3. Fig. 7 and Fig. 6 visualize our predicted depth and recovered point clouds. The results are not only clean, accurate, and complete, but also preserve fine-grained geometric details, clearly demonstrating a superiority over other methods. Even more notably, our much smaller DA3-Large (0.30B parameters) demonstrates remarkable efficiency. Despite being $3\times$ smaller, it surpasses the prior SOTA VGGT (1.19B parameters) in five out of the ten settings, with particularly strong performance on the ETH3D.

When camera poses are available, both our method and MapAnything can exploit them for improved results, and other methods also benefit from ground-truth pose fusion. Our model shows clear gains on most datasets except 7Scenes, where the limited video setting already saturates performance and reduces the benefit of pose conditioning. Notably, with pose conditioning, performance gains from scaling model size are smaller than in pose-free models, **indicating that pose estimation scales more strongly than depth estimation and requires larger models to fully realize improvements.**

Monocular depth accuracy also reflects geometry quality. As shown in Tab. 4, on the standard monocular depth benchmarks reported in [113], our model outperforms VGGT and Depth Anything 2. For reference, we also include the results of our teacher model.

Table 4 Monocular depth comparisons. $\delta_1 \uparrow$

Method	KITTI	NYU	SINTEL	ETH3D	DIODE	Rank
DA2	94.6	97.9	77.2	86.5	95.2	2.60
VGGT	91.7	97.9	67.9	97.5	95.3	3.75
DA3	95.3	97.4	75.5	98.6	95.4	2.20
Teacher	97.2	97.9	81.4	99.8	96.6	1.00

Visual rendering. To fairly evaluate feed-forward novel view synthesis (FF-NVS), we compare against three recent 3DGS models—pixelSplat [11], MVSplat [13], and DepthSplat [108]—and further test alternative frameworks by replacing our geometry backbone with Fast3R [111], MV-DUSt3R [85], and VGGT [91]. All models are trained on DL3DV-10K training set under a unified protocol and evaluated on our benchmark (Sec. 6.3).

As reported in Tab. 5, all models perform substantially better on DL3DV than on the other datasets, suggesting that 3DGS-based NVS is sensitive to trajectory and pose distributions standardized by DL3DV, rather than scene content. Comparing the two groups, geometry-model-based frameworks consistently outperform specialized feed-forward models, demonstrating that a simple backbone plus DPT head can




Figure 6 Comparisons of point cloud quality. Our model produces point clouds that are more geometrically regular and substantially less noisy than those generated by other methods.




Figure 7 Comparisons of depth quality. Compared with other methods, our depth maps exhibit finer structural detail and higher semantic correctness across diverse scenes.

surpass complex task-specific designs. The advantage stems from large-scale pretraining, which enables better generalization and scalability than approaches relying on epipolar transformers, cost volumes, or cascaded modules. Within this group, NVS performance correlates with geometry estimation capability, making DA3 the strongest backbone. Looking forward, we expect FF-NVS can be effectively addressed with simple architectures leveraging pretrained geometry backbones, and that the strong spatial understanding of DA3 will benefit other 3D vision tasks.

7.2 Analysis for Depth Anything 3

Training our DA3-Giant model requires $128 \times H100$ GPUs for approximately 10 days. To reduce carbon footprint and computational cost, all ablation experiments reported in this section are conducted using the ViT-L backbone with a maximum of 10 views, requiring approximately 4 days on $32 \times H100$ GPUs.

7.2.1 Sufficiency of the Depth-Ray Representation

To validate our depth-ray representation, we compare different prediction combinations summarized in Tab. 6. All models use a ViT-L backbone, identical training settings (view size: 10, batch size: 128, steps: 120k). We evaluate four heads: 1) **depth** for dense depth maps; 2) **pcd** for direct 3D point clouds; 3) **cam** for 9-DoF camera pose $\mathbf{c} = (\mathbf{t}, \mathbf{q}, \mathbf{f})$; and 4) our proposed **ray**, predicting per-pixel ray maps (Sec. 3.1). The **ray** head uses a Dual-DPT architecture, while **pcd** uses a separate DPT head. For models without **pcd**, point clouds

Table 5 Comparisons with SOTA methods on NVS task. We report NVS comparisons with existing feed-forward 3DGS models and counterparts using other backbones. For each scene, we use 12 input context views and test on target views sampled every 8 views over a set of over 300 views. Image resolution is 270×480 .

Methods	In-domain Dataset			Out-of-domain Datasets					
	DL3DV-Benchmarks (140)			Tanks and Temples (6)			MegaDepth (19)		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
pixelSplat	16.55	0.456	0.480	13.81	0.347	0.558	13.87	0.367	0.561
MVSplat	18.13	0.559	0.393	14.81	0.391	0.508	14.67	0.398	0.533
DepthSplat	19.24	0.620	0.322	15.80	0.474	0.418	15.90	0.471	0.450
Fast3R	19.30	0.604	0.320	16.24	0.478	0.409	16.43	0.493	0.421
MV-DUST3R	20.01	0.645	0.294	17.04	0.529	0.370	16.20	0.484	0.437
VGGT	20.96	0.697	0.253	17.18	0.550	0.347	16.45	0.500	0.417
DAv3 (Ours)	21.33	0.711	0.241	18.10	0.578	0.311	17.89	0.561	0.351

Table 6 Ablations of prediction-target combinations. Note that all experiments in this table do not have camera condition token. The **best** and second best are highlighted.

Methods	HiRoom		ETH3D		DTU		7Scenes		ScanNet++	
	Auc3 $\uparrow$	F1 $\uparrow$	Auc3 $\uparrow$	F1 $\uparrow$	Auc3 $\uparrow$	CD $\downarrow$	Auc3 $\uparrow$	F1 $\uparrow$	Auc3 $\uparrow$	F1 $\uparrow$
depth + pcd + cam	9.1	12.8	19.0	<u>60.4</u>	42.3	4.918	20.8	43.4	22.0	43.0
depth + cam	10.8	16.5	9.9	48.0	23.3	5.316	13.0	38.5	13.3	41.0
depth + ray	48.7	60.3	25.5	65.4	<u>46.5</u>	<u>3.919</u>	<u>24.0</u>	46.5	35.5	<u>53.4</u>
depth + ray + cam	<u>37.2</u>	<u>45.4</u>	<u>22.3</u>	59.4	56.3	3.066	25.7	<u>45.6</u>	<u>34.1</u>	56.5

are obtained by combining **depth** with camera parameters from **ray** or **cam**. As shown in Table 6, the minimal **depth + ray** configuration consistently outperforms **depth + pcd + cam** and **depth + cam** across all datasets and metrics, achieving nearly 100% relative gain in Auc3 over **depth + cam**. Adding an auxiliary **cam** head (**depth + ray + cam**) yields no further benefit, confirming the sufficiency of the depth-ray representation. We adopt **depth + ray + cam** as our final representation, as the camera head incurs negligible computational overhead, amounting to approximately 0.1% of the computation cost of the main backbone.

7.2.2 Sufficiency of a Single Plain Transformer

We compare a standard ViT-L backbone with a VGGT-style architecture that stacks two distinct transformers, tripling the block count. For fair capacity comparison, the VGGT-style model uses smaller ViT-B backbones, yielding a similar parameter size to our ViT-L. Our backbone supports two attention strategies: **Full Alt.**, which alternates cross-view/within-view attention in all layers ($L = L_g$), and our default partial alternation. As shown in Table 7, the VGGT-style model drops to 79.8% of our baseline performance, confirming the superiority of a single-transformer design at similar scale. We attribute this gap to full pretraining of our backbone versus two-thirds untrained blocks in VGGT. Moreover, the **Full Alt.** variant degrades across nearly all metrics—except F1 on 7Scenes—indicating that partial alternation is the more effective and robust strategy.

7.2.3 Ablation and Analysis

Dual-DPT Head. We assess the effectiveness of the dual-DPT head via an ablation in which two separate DPT heads predict depth and ray maps independently. Results are reported in Tab. 7, item (d). Compared with the model equipped with the dual-DPT head, the variant without it shows consistent drops across metrics, confirming the effectiveness of our dual-DPT design.

Table 7 Ablation study. We evaluate three architectural designs with comparable model sizes (a-c), the effects of the dual-DPT head (d), teacher label supervision (e), and the pose conditioning module (f-g). The **best** and second best are highlighted. Methods marked with "*" are evaluated with ground-truth pose fusion.

Methods	HiRoom		ETH3D		DTU		7Scenes		ScanNet++	
	Auc3↑	F1↑	Auc3↑	F1↑	Auc3↑	CD↓	Auc3↑	F1↑	Auc3↑	F1↑
a. Proposed Arch.	39.2	47.0	21.0	55.4	45.8	3.82	26.2	47.6	30.3	51.1
b. VGGT Style	3.72	14.5	2.31	27.4	1.38	6.93	0.97	21.4	2.03	12.2
c. Full Alt.	<u>24.7</u>	<u>29.3</u>	<u>13.1</u>	<u>51.9</u>	<u>44.6</u>	<u>4.23</u>	<u>21.1</u>	48.6	<u>27.7</u>	<u>47.5</u>
d. w/o Dual DPT	5.59	11.5	13.6	33.4	21.7	5.14	14.2	49.4	26.5	46.6
e. w/o Teacher	11.2	16.0	16.2	57.6	52.5	3.29	23.3	40.3	26.2	47.7
f. w/o Pose Cond.*		65.8		63.2		3.65		58.4		62.8
g. w/ Pose Cond.*		73.8		70.9		2.14		46.0		65.7

Table 8 Comparison of Models with Parameters and Running Speed. The maximum number of images was tested on an 80 GB A100 GPU. If we store some intermediate tokens in CPU memory, we could process many more images. The running speed was measured on an A100 GPU with a scene of 32 images, and we report the average speed per image. The image resolution is 504×336 .

Model	Max # of Images	Parameters			Running Speed
		Backbone	DualDPT	CameraHead	
VGGT(Reference)	400-500	0.91B	0.064B	0.22B	34.1 FPS
DA3-Giant	900-1000	1.130B	0.050B	0.018B	37.6 FPS
DA3-Large	1500-1600	0.300B	0.047B	0.008B	78.37 FPS
DA3-Base	2100-2200	0.086B	0.015B	0.004B	126.5 FPS
DA3-Small	4000-4100	0.022B	0.003B	0.001B	160.5 FPS

Teacher model supervision. We ablate the use of teacher model labels as supervision, with quantitative results reported in Tab. 7, item (e). Training without teacher labels yields a slight improvement on DTU but leads to performance drops on 7Scenes and ScanNet++. Notably, the degradation is pronounced on HiRoom. We attribute this to HiRoom’s synthetic nature and its ground truth containing abundant fine structures; supervision from the teacher helps the student capture such details more accurately. Qualitative comparisons in Fig. 8 corroborate this trend: models trained with teacher-label supervision produce depth maps with substantially richer detail and finer structures.

Pose conditioning. To assess the pose-conditioning module, we ablate it on the ViT-L backbone and report results in Tab. 7, items (f) and (g). Unlike other entries in the table, these two are evaluated with ground-truth pose fusion (marked with “*”), whereas the rest use predicted pose fusion. Across metrics, configurations with pose conditioning consistently outperform those without, confirming the effectiveness of the pose-conditioning module.

Running time. We present analysis on Parameters, max number of images and running speed in Tab. 8

More visualizations. We provide additional visualizations of camera pose and depth estimation on in-the-wild scenes in Fig. 9, demonstrating the robustness and quality of our model across diverse real-world scenarios.

7.3 Analysis for Depth-Anything-3-Monocular



Figure 8 Comparison of teacher-label supervision. Supervision with teacher-generated labels yields depth maps with substantially richer detail and finer structures.

Table 9 Ablations for teacher model. Training with V3 datasets and multi-resolution strategy yields the best performance. Depth-based geometry achieves the best AbsRel and SqRel. The full teacher-loss outperforms other variants. (AbsRel: ↓, SqRel: ↓, δ_1 : ↑). The results are averaged over KITTI, NYU, ETH3D, SUN-RGBD and DIODE.

Data				Geometry				Loss			
Data	δ_1	AbsRel	SqRel	Geometry	δ_1	AbsRel	SqRel	Loss	δ_1	AbsRel	SqRel
V2	0.919	0.087	0.596	Disparity	0.919	<u>0.095</u>	1.033	MAE-Loss	0.918	0.089	0.637
V3	<u>0.929</u>	<u>0.079</u>	<u>0.508</u>	Pointmap	0.912	0.096	<u>0.693</u>	w/o Dist. Nor.	0.918	0.087	<u>0.600</u>
V3 + mr.	0.938	0.072	0.452	Depth	<u>0.918</u>	0.089	0.637	Full loss	0.919	0.087	0.596

7.3.1 Teacher Model

The teacher model’s metrics are reported in Tab. 4. Our new teacher consistently outperforms DA2 across all datasets, with the sole exception of NYU, where performance is on par with DA2. For the teacher ablation, we employ a ViT-L backbone and a batch size of 64. Evaluation follows the DA2 benchmark protocol, and we additionally report Squared Relative Error (SqRel), defined as the mean squared error between predictions and ground truth normalized by the ground truth. As shown in Tab. 9, across geometries, depth emerges as the most effective target compared with disparity and point maps. For training objectives, the full teacher loss proposed in this work outperforms both the DA2 loss and a variant without proposed normal-loss term. Finally, data scaling contribute notably to performance: upgrading datasets from V2 to V3 and adopting a multi-resolution training strategy yield consistent improvements in the teacher’s final metrics.

7.3.2 Student Model

As shown in Tab. 10, our monocular student model with a ViT-L backbone outperforms the DA2 student across all evaluation datasets. Notably, on the ETH3D [72] benchmark the new monocular student achieves an improvement of over 10% compared with DA2. The improved performance is attributed to the enhanced teacher model with better geometry supervision and the scaled training data (V3). On challenging datasets like SINTEL, our student also demonstrates substantial gains (+5.1%), validating the effectiveness of our teacher-student distillation framework.

7.4 Analysis for Depth-Anything-3-Metric

We compare with state-of-the-art metric depth estimation methods, including DepthPro [7], Metric3D v2 [37], UniDepthv1 [65] and UniDepthv2 [66], on 5 benchmarks: NYUv2 [76], KITTI [25], ETH3D [72],




Figure 9 Visualizations of camera pose and depth estimation on in-the-wild scenes.

SUN-RGBD [81] and DIODE (indoor) [88].

As shown in Tab. 11, DA3-metric achieves **state-of-the-art performance on ETH3D** ($\delta_1 = 0.917$, AbsRel = 0.104), substantially outperforming the second-best method UniDepthv2 ($\delta_1 = 0.863$) by a large margin. DA3-metric also achieves **best performance on SUN-RGBD** for AbsRel (0.105) and **second-best on DIODE** ($\delta_1 = 0.838$, AbsRel = 0.128). While UniDepthv1 and UniDepthv2 achieve the best results on NYUv2 and KITTI, DA3-metric demonstrates strong generalization and competitive performance across all benchmarks, particularly excelling on diverse outdoor scenes like ETH3D.

We ablate the Teacher supervision in Tab. 11. The results show interesting trade-offs: removing Teacher supervision slightly improves metrics on NYUv2 and KITTI, while maintaining comparable performance on other datasets. As shown in Fig. 10, Teacher supervision significantly improves sharpness and fine detail quality, demonstrating that Teacher provides complementary knowledge beyond standard metrics.

7.5 Analysis for Feed-forward 3DGS

We retrain all compared feed-forward 3DGS models, ensuring that the training configuration matches the testing setup by using 12 input context views selected through farthest point sampling. We apply engineering optimizations such as flash attention and fully shared data parallelism to enable all models to process 12

Table 10 Monocular student depth comparisons. $\delta_1 \uparrow$

Method	KITTI	NYU	SINTEL	ETH3D	DIODE
DA2	94.6	97.9	77.2	86.5	95.2
mono-student	97.1	98.0	82.3	98.8	96.5

Table 11 Comparison with state-of-the-arts on metric depth estimation. The **best** and **second best** are highlighted. Bottom rows show ablation results with and without teacher supervision. Note that the ablation setting is slightly different from the final model on training resolution, which leads to minor differences in performance.

Methods	NYUv2		KITTI		ETH3D		SUN-RGBD		DIODE	
	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$
DepthPro [7]	0.932	0.093	0.843	0.121	0.386	0.349	0.950	0.126	0.734	0.173
Metric3D v2 [36]	<u>0.971</u>	0.067	<u>0.976</u>	0.051	0.830	<u>0.138</u>	0.954	0.132	0.018	0.154
UniDepthv1 [65]	0.980	0.061	0.978	0.051	0.234	0.464	<u>0.971</u>	0.113	0.570	0.266
UniDepthv2 [66]	0.968	<u>0.064</u>	0.968	0.076	<u>0.863</u>	0.152	0.977	<u>0.111</u>	0.856	0.123
DA3-metric	0.963	0.070	0.953	0.086	0.917	0.104	<u>0.973</u>	0.105	<u>0.838</u>	<u>0.128</u>
w/ teacher	0.966	0.073	0.947	0.086	0.906	0.105	0.973	0.104	0.824	0.132
w/o teacher	0.969	0.066	0.965	0.067	0.907	0.105	0.975	0.099	0.816	0.134




Figure 10 Effectiveness of Teacher model for supervising metric depth estimation. Incorporating Teacher model for supervision significantly improves the metric depth sharpness.

input views efficiently. Depth training loss are incorporated for all baselines to ensure stable training and convergence. All models are trained on 8 A100 GPUs for 200K steps with a batch size of 1, except for pixelSplat, which is trained for 100K steps due to rather slow epipolar attention. All results are reported at $H \times W = 270 \times 480$.

Visual quality analysis. We present visual comparisons with other models in Fig. 11 under novel view synthesis settings. As illustrated, simply augmenting our DA3 model with a 3D Gaussian DPT head yields significantly improved rendering quality over existing state-of-the-art approaches. Our model demonstrates particular strength in challenging regions, such as thin structures (*e.g.*, columns in the first and third scenes) and large-scale outdoor environments with wide-baseline input views (last two scenes), as shown in Fig. 11. These results underscore the importance of a robust geometry backbone for high-quality visual rendering, consistent with our quantitative findings in Tab. 5. We anticipate that the strong geometric understanding of DA3 will also benefit other 3D vision tasks.

8 Conclusion and Discussion

Depth Anything 3 shows that a plain transformer, trained on depth-and-ray targets with teacher-student supervision, can unify any-view geometry without ornate architectures. Scale-aware depth, per-pixel rays, and adaptive cross-view attention let the model inherit strong pretrained features while remaining lightweight and easy to extend. On the proposed visual geometry benchmark the approach sets new pose and reconstruction records, with both giant and compact variants surpassing prior models, while the same backbone powers efficient feed-forward novel view synthesis model.

We view Depth Anything 3 as a step toward versatile 3D foundation models. Future work can extend its reasoning to dynamic scenes, integrate language and interaction cues, and explore larger-scale pretraining to close the loop between geometry understanding and actionable world models. We hope the model and dataset releases, benchmark, and simple modeling principles offered here catalyze broader research on general-purpose 3D perception.

Acknowledgement

We thank Xiaowei Zhou, Sida Peng and Hengkai Guo for their valuable discussions during the development of this project. We are also grateful to Yang Zhao for his engineering support. The input images in the teaser demo were extracted from a publicly available YouTube video [21], credited to the original creator.




Figure 11 Qualitative comparisons with state-of-the-art methods for visual rendering. The first column shows the selected input views, while the remaining columns display novel views rendered by comparison models and ground truth. For each scene, two rendered novel viewpoints are presented in consecutive rows. The first three scenes are from DL3DV, the following two are from Tanks and Temples, and the last three are from MegaDepth. Compared to other methods, our model consistently achieves superior rendering quality across diverse and challenging scenes.

References

- [1] Henrik Aanaes, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjorholm Dahl. Large-scale data for multiple-view stereopsis. *Int. J. Comput. Vis.*, 120(2):153–168, 2016.
- [2] Yousset I Abdel-Aziz, Hauck Michael Karara, and Michael Hauck. Direct linear transformation from comparator coordinates into object space coordinates in close-range photogrammetry. *Photogrammetric engineering & remote sensing*, 81(2):103–107, 2015.
- [3] Eduardo Arnold, Jamie Wynn, Sara Vicente, Guillermo Garcia-Hernando, Aron Monszpart, Victor Prisacariu, Daniyar Turmukhambetov, and Eric Brachmann. Map-free visual relocalization: Metric pose relative to a single image. In *European Conference on Computer Vision*, pages 690–708. Springer, 2022.
- [4] Martha E Arterberry and Albert Yonas. Perception of three-dimensional shape specified by optic flow by 8-week-old infants. *Perception & Psychophysics*, 62(3):550–556, 2000.
- [5] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *Adv. Neural Inform. Process. Syst.*, 2021.
- [6] Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. *arXiv preprint arXiv:2410.02073*, 2024.
- [7] Alexey Bochkovskiy, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. In *Int. Conf. Learn. Represent.*, 2025.
- [8] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432, 2001.
- [9] Yann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2, 2020.
- [10] Yann Cabon, Lucas Stoffl, Leonid Antsfeld, Gabriela Csurka, Boris Chidlovskii, Jerome Revaud, and Vincent Leroy. Must3r: Multi-view network for stereo 3d reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1050–1060, 2025.
- [11] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 19457–19467, 2024.
- [12] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 14124–14133, 2021.
- [13] Yuedong Chen, Haoqi Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *Eur. Conf. Comput. Vis.*, pages 370–386. Springer, 2024.
- [14] Yuedong Chen, Haoqi Xu, Qianyi Wu, Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. Explicit correspondence matching for generalizable neural radiance fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2025.
- [15] Jaehoon Cho, Dongbo Min, Youngjung Kim, and Kwanghoon Sohn. Diml/cvl rgb-d dataset: 2m rgb-d images of natural indoor and outdoor scenes. *arXiv preprint arXiv:2110.11590*, 2021.
- [16] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [17] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13142–13153, 2023.
- [18] Junyuan Deng, Heng Li, Tao Xie, Weiqiang Ren, Qian Zhang, Ping Tan, and Xiaoyang Guo. Sail-recon: Large sfm by augmenting scene regression with localization. *arXiv preprint arXiv:2508.17972*, 2025.

- [19] Kai Deng, Zexin Ti, Jiawei Xu, Jian Yang, and Jin Xie. Vggt-long: Chunk it, loop it, align it – pushing vggt’s limits on kilometer-scale long rgb sequences. *arXiv preprint arXiv:2507.11539*, 2025.
- [20] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *CVPR workshops*, pages 224–236, 2018.
- [21] MTS Drones. Drone australia gliding ep025: Sydney views | opera house, harbour bridge & hyde park | dji mavic 4k. <https://www.youtube.com/watch?v=qbgKDaGraTA>, 2024. Accessed: Sep. 25, 2025. Used under YouTube Standard License.
- [22] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8092–8101, 2019.
- [23] David Eigen, Christian Puhersch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Adv. Neural Inform. Process. Syst.*, 2014.
- [24] Mathias Gehrig, Willem Aarents, Daniel Gehrig, and Davide Scaramuzza. Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters*, 6(3):4947–4954, 2021.
- [25] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11):1231–1237, 2013.
- [26] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11):1231–1237, 2013.
- [27] Yotam Gil, Shay Elmaleh, Harel Haim, Emanuel Marom, and Raja Giryes. Online training of stereo self-calibration using monocular depth estimation. *IEEE Transactions on Computational Imaging*, 7:812–823, 2021.
- [28] Shrisudhan Govindarajan, Daniel Rebain, Kwang Moo Yi, and Andrea Tagliasacchi. Radiant foam: Real-time differentiable ray tracing. In *Int. Conf. Comput. Vis.*, 2025.
- [29] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5354–5363, 2024.
- [30] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020.
- [31] Haoyu Guo, He Zhu, Sida Peng, Haotong Lin, Yunzhi Yan, Tao Xie, Wenguan Wang, Xiaowei Zhou, and Hujun Bao. Multi-view reconstruction via sfm-guided monocular depth estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5272–5282, 2025.
- [32] Jose L. Gómez, Manuel Silva, Antonio Seoane, Agnès Borràs, Mario Noriega, German Ros, Jose A. Iglesias-Guitian, and Antonio M. López. All for one, and one for all: Urbansyn dataset, the third musketeer of synthetic driving scenes. *Neurocomputing*, 637:130038, 2025. ISSN 0925-2312. doi: <https://doi.org/10.1016/j.neucom.2025.130038>. URL <https://www.sciencedirect.com/science/article/pii/S0925231225007106>.
- [33] Xingyi He, Jiaming Sun, Yifan Wang, Sida Peng, Qixing Huang, Hujun Bao, and Xiaowei Zhou. Detector-free structure from motion. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21594–21603, 2024.
- [34] Benno Heigl, Reinhard Koch, Marc Pollefeys, Joachim Denzler, and Luc Van Gool. Plenoptic modeling and rendering from image sequences taken by a hand-held camera. In *Mustererkennung 1999: 21. DAGM-Symposium Bonn, 15.–17. September 1999*, pages 94–101. Springer, 1999.
- [35] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. In *Int. Conf. Learn. Represent.*, 2024.
- [36] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [37] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *TPAMI*, 2024.

- [38] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018.
- [39] Wonbong Jang, Philippe Weinzaepfel, Vincent Leroy, Lourdes Agapito, and Jerome Revaud. Pow3r: Empowering unconstrained 3d reconstruction with camera and scene priors. In IEEE Conf. Comput. Vis. Pattern Recog., pages 1071–1081, 2025.
- [40] Hanwen Jiang, Zexiang Xu, Desai Xie, Ziwen Chen, Haian Jin, Fujun Luan, Zhixin Shu, Kai Zhang, Sai Bi, Xin Sun, et al. Megasynt: Scaling up 3d scene reconstruction with synthesized data. In Proceedings of the Computer Vision and Pattern Recognition Conference, pages 16441–16452, 2025.
- [41] Lihan Jiang, Yucheng Mao, Linning Xu, Tao Lu, Kerui Ren, Yichen Jin, Xudong Xu, Mulin Yu, Jiangmiao Pang, Feng Zhao, et al. Anysplat: Feed-forward 3d gaussian splatting from unconstrained views. ACM Trans. Graph., 2025.
- [42] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In CVPR, pages 9492–9502, 2024.
- [43] Nikhil Keetha, Norman Müller, Johannes Schönberger, Lorenzo Porzi, Yuchen Zhang, Tobias Fischer, Arno Knapitsch, Duncan Zauss, Ethan Weber, Nelson Antunes, Jonathon Luiten, Manuel Lopez-Antequera, Samuel Rota Bulò, Christian Richardt, Deva Ramanan, Sebastian Scherer, and Peter Kotschieder. MapAnything: Universal feed-forward metric 3D reconstruction, 2025. arXiv preprint arXiv:2509.13414.
- [44] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph., 42(4):139–1, 2023.
- [45] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Trans. Graph., 36(4):1–13, 2017.
- [46] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Trans. Graph., 36(4):1–13, 2017.
- [47] Hoang-An Le, Thomas Mensink, Partha Das, Sezer Karaoglu, and Theo Gevers. Eden: Multimodal synthetic dataset of enclosed garden scenes. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 1579–1589, 2021.
- [48] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In Eur. Conf. Comput. Vis., pages 71–91. Springer, 2024.
- [49] Marc Levoy and Pat Hanrahan. Light field rendering. ACM Trans. Graph., 1996.
- [50] Yixuan Li, Lihan Jiang, Linning Xu, Yuanbo Xiangli, Zhenzhi Wang, Dahua Lin, and Bo Dai. Matrixcity: A large-scale city dataset for city-scale neural rendering and beyond. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3205–3215, 2023.
- [51] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In IEEE Conf. Comput. Vis. Pattern Recog., pages 2041–2050, 2018.
- [52] Haotong Lin, Sida Peng, Zhen Xu, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Efficient neural radiance fields for interactive free-viewpoint video. In SIGGRAPH Asia 2022 Conference Papers, pages 1–9, 2022.
- [53] Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. D13dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In IEEE Conf. Comput. Vis. Pattern Recog., pages 22160–22169, 2024.
- [54] Dominic Maggio, Hyungtae Lim, and Luca Carlone. Vggt-slam: Dense rgb slam optimized on the sl(4) manifold. arXiv preprint arXiv:2505.12549, 2025.
- [55] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J Davison. Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation? In Proceedings of the IEEE International Conference on Computer Vision, pages 2678–2687, 2017.
- [56] Lukas Mehl, Jenny Schmalfuss, Azin Jahedi, Yaroslava Nalivayko, and Andrés Bruhn. Spring: A high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4981–4991, 2023.

- [57] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In Eur. Conf. Comput. Vis., 2020.
- [58] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: A versatile and accurate monocular slam system. IEEE transactions on robotics, 31(5):1147–1163, 2015.
- [59] Riku Murai, Eric Dexheimer, and Andrew J Davison. Mast3r-slam: Real-time dense slam with 3d reconstruction priors. In IEEE Conf. Comput. Vis. Pattern Recog., pages 16695–16705, 2025.
- [60] Simon Niklaus, Long Mai, Jimei Yang, and Feng Liu. 3d ken burns effect from a single image. ACM Transactions on Graphics (ToG), 38(6):1–15, 2019.
- [61] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193, 2023.
- [62] Linfei Pan, Dániel Baráth, Marc Pollefeys, and Johannes L Schönberger. Global structure-from-motion revisited. In Eur. Conf. Comput. Vis., pages 58–77. Springer, 2024.
- [63] Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Yuheng Carl Ren. Aria digital twin: A new benchmark dataset for egocentric 3d machine perception. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 20133–20143, 2023.
- [64] William Peebles and Saining Xie. Scalable diffusion models with transformers. In Proceedings of the IEEE/CVF international conference on computer vision, pages 4195–4205, 2023.
- [65] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. In IEEE Conf. Comput. Vis. Pattern Recog., pages 10106–10116, 2024.
- [66] Luigi Piccinelli, Christos Sakaridis, Yung-Hsu Yang, Mattia Segu, Siyuan Li, Wim Abbeloos, and Luc Van Gool. Unidepthv2: Universal monocular metric depth estimation made simpler. arXiv preprint arXiv:2502.20110, 2025.
- [67] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In Int. Conf. Comput. Vis., pages 12179–12188, 2021.
- [68] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In Int. Conf. Comput. Vis., pages 10901–10911, 2021.
- [69] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In Proceedings of the IEEE/CVF international conference on computer vision, pages 10912–10922, 2021.
- [70] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In IEEE Conf. Comput. Vis. Pattern Recog., 2016.
- [71] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In Eur. Conf. Comput. Vis., 2016.
- [72] 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 multi-camera videos. In IEEE Conf. Comput. Vis. Pattern Recog., pages 3260–3269, 2017.
- [73] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In IEEE Conf. Comput. Vis. Pattern Recog., volume 1, pages 519–528. IEEE, 2006.
- [74] 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 IEEE Conf. Comput. Vis. Pattern Recog., pages 2930–2937, 2013.
- [75] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In Eur. Conf. Comput. Vis., pages 746–760. Springer, 2012.

- [76] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *Eur. Conf. Comput. Vis.*, pages 746–760. Springer, 2012.
- [77] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Adv. Neural Inform. Process. Syst.*, 32, 2019.
- [78] Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Adv. Neural Inform. Process. Syst.*, 34:19313–19325, 2021.
- [79] Brandon Smart, Chuanxia Zheng, Iro Laina, and Victor Adrian Prisacariu. Splatt3r: Zero-shot gaussian splatting from uncalibrated image pairs. *arXiv preprint arXiv:2408.13912*, 2024.
- [80] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. *ACM Trans. Graph.*, pages 835–846, 2006.
- [81] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 567–576, 2015.
- [82] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.
- [83] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.
- [84] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [85] Zhenggang Tang, Yuchen Fan, Dilin Wang, Hongyu Xu, Rakesh Ranjan, Alexander Schwing, and Zhicheng Yan. Mv-dust3r+: Single-stage scene reconstruction from sparse views in 2 seconds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5283–5293, 2025.
- [86] Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. *arXiv preprint arXiv:1812.04605*, 2018.
- [87] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(4):376–380, 2002.
- [88] Igor Vasiljevic, Nick Kolkin, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z Dai, Andrea F Daniele, Mohammadreza Mostajabi, Steven Basart, Matthew R Walter, et al. Diode: A dense indoor and outdoor depth dataset. *arXiv preprint arXiv:1908.00463*, 2019.
- [89] Jianyuan Wang, Christian Rupprecht, and David Novotny. Posediffusion: Solving pose estimation via diffusion-aided bundle adjustment. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9773–9783, 2023.
- [90] Jianyuan Wang, Nikita Karaev, Christian Rupprecht, and David Novotny. Vggsfm: Visual geometry grounded deep structure from motion. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21686–21697, 2024.
- [91] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5294–5306, 2025.
- [92] Kaixuan Wang and Shaojie Shen. Flow-motion and depth network for monocular stereo and beyond. *IEEE Robotics and Automation Letters*, 5(2):3307–3314, 2020.
- [93] Qiang Wang, Shizhen Zheng, Qingsong Yan, Fei Deng, Kaiyong Zhao, and Xiaowen Chu. Irs: A large naturalistic indoor robotics stereo dataset to train deep models for disparity and surface normal estimation. *arXiv preprint arXiv:1912.09678*, 2019.

- [94] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10510–10522, 2025.
- [95] Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. Moge: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5261–5271, 2025.
- [96] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 20697–20709, 2024.
- [97] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 20697–20709, 2024.
- [98] Wenshan Wang, DeLong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. Tartanair: A dataset to push the limits of visual slam. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4909–4916. IEEE, 2020.
- [99] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. π^3 : Scalable permutation-equivariant visual geometry learning, 2025. URL <https://arxiv.org/abs/2507.13347>.
- [100] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. Foundationstereo: Zero-shot stereo matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5249–5260, 2025.
- [101] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Adv. Neural Inform. Process. Syst.*, 2021.
- [102] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Liang Pan Jiawei Ren, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [103] Hongchi Xia, Yang Fu, Sifei Liu, and Xiaolong Wang. Rgb-d objects in the wild: Scaling real-world 3d object learning from rgb-d videos, 2024.
- [104] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*, 2024.
- [105] Pengchuan Xiao, Zhenlei Shao, Steven Hao, Zishuo Zhang, Xiaolin Chai, Judy Jiao, Zesong Li, Jian Wu, Kai Sun, Kun Jiang, et al. Pandaset: Advanced sensor suite dataset for autonomous driving. In *2021 IEEE international intelligent transportation systems conference (ITSC)*, pages 3095–3101. IEEE, 2021.
- [106] Gangwei Xu, Xianqi Wang, Xiaohuan Ding, and Xin Yang. Iterative geometry encoding volume for stereo matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21919–21928, 2023.
- [107] Haoifei Xu, Anpei Chen, Yuedong Chen, Christos Sakaridis, Yulun Zhang, Marc Pollefeys, Andreas Geiger, and Fisher Yu. Murf: multi-baseline radiance fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 20041–20050, 2024.
- [108] Haoifei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. Depthsplat: Connecting gaussian splatting and depth. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16453–16463, 2025.
- [109] Guorun Yang, Xiao Song, Chaoqin Huang, Zhidong Deng, Jianping Shi, and Bolei Zhou. Drivingstereo: A large-scale dataset for stereo matching in autonomous driving scenarios. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [110] 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. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21924–21935, 2025.
- [111] 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. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21924–21935, 2025.

- [112] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *Adv. Neural Inform. Process. Syst.*, 37:21875–21911, 2024.
- [113] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. In *Adv. Neural Inform. Process. Syst.*, 2024.
- [114] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Eur. Conf. Comput. Vis.*, pages 767–783, 2018.
- [115] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1790–1799, 2020.
- [116] Botao Ye, Sifei Liu, Haoifei Xu, Xueting Li, Marc Pollefeys, Ming-Hsuan Yang, and Songyou Peng. No pose, no problem: Surprisingly simple 3d gaussian splats from sparse unposed images. In *Int. Conf. Learn. Represent.*, 2024.
- [117] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Int. Conf. Comput. Vis.*, pages 12–22, 2023.
- [118] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *CVPR*, pages 9043–9053, 2023.
- [119] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4578–4587, 2021.
- [120] Amir R Zamir, Alexander Sax, , William B Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *IEEE Conf. Comput. Vis. Pattern Recog.* IEEE, 2018.
- [121] 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. In *Int. Conf. Learn. Represent.*, 2025.
- [122] Shangzhan Zhang, Jianyuan Wang, Yinghao Xu, Nan Xue, Christian Rupprecht, Xiaowei Zhou, Yujun Shen, and Gordon Wetzstein. Flare: Feed-forward geometry, appearance and camera estimation from uncalibrated sparse views. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21936–21947, 2025.
- [123] Yi Zhang, Weichao Qiu, Qi Chen, Xiaolin Hu, and Alan Yuille. Unrealstereo: Controlling hazardous factors to analyze stereo vision. In *2018 International Conference on 3D Vision (3DV)*, pages 228–237. IEEE, 2018.
- [124] Zhe Zhang, Rui Peng, Yuxi Hu, and Ronggang Wang. Geomvsnet: Learning multi-view stereo with geometry perception. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21508–21518, 2023.
- [125] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. In *European Conference on Computer Vision*, pages 519–535. Springer, 2020.
- [126] Peng Zheng, Dehong Gao, Deng-Ping Fan, Li Liu, Jorma Laaksonen, Wanli Ouyang, and Nicu Sebe. Bilateral reference for high-resolution dichotomous image segmentation. *CAAI Artificial Intelligence Research*, 2024.
- [127] Yang Zheng, Adam W Harley, Bokui Shen, Gordon Wetzstein, and Leonidas J Guibas. Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19855–19865, 2023.
- [128] Yang Zhou, Yifan Wang, Jianjun Zhou, Wenzheng Chang, Haoyu Guo, Zizun Li, Kaijing Ma, Xinyue Li, Yating Wang, Haoyi Zhu, Mingyu Liu, Dingning Liu, Jiange Yang, Zhoujie Fu, Junyi Chen, Chunhua Shen, Jiangmiao Pang, Kaipeng Zhang, and Tong He. Omniworld: A multi-domain and multi-modal dataset for 4d world modeling, 2025. URL <https://arxiv.org/abs/2509.12201>.
- [129] Rui Zhu, Second Author, and Third Author. Svlightverse: Large-scale photorealistic indoor dataset with spatially-varying hdri lighting. <https://jerrypiglet.github.io/SVLightVerse/>, 2025. Project page. Affiliations: University of California San Diego; PICO (ByteDance); KooLab (Manycore); Rembrand. Accessed: 2025-11-11.
- [130] Zihan Zhu, Songyou Peng, Viktor Larsson, Zhaopeng Cui, Martin R Oswald, Andreas Geiger, and Marc Pollefeys. Nicer-slam: Neural implicit scene encoding for rgb slam. In *3DV*, pages 42–52. IEEE, 2024.

Appendix

While synthetic datasets provide large-scale training data with ground-truth depth annotations, many contain quality issues such as invalid backgrounds, spatial misalignments, clipping artifacts, and erroneous depth values that can degrade model training. We therefore apply careful preprocessing to filter problematic samples and clip unrealistic depth ranges, ensuring high-quality supervision for our teacher model.

A Data Processing

We preprocess the following raw training datasets as follows to train the teacher model.

TartanAir. We remove the `amusement` scene from training due to its invalid background (skybox) (Fig. 12). We clip the maximum depth values of `carwelding`, `hospital`, `ocean`, `office` and `office2` scenes at 80, 30, 1000, 30 and 30, respectively.



Figure 12 Invalid background of `amusement` scene in TartanAir dataset.

IRS. We noticed that some of the scenes in IRS exhibit spatial misalignment between the image and depth maps (Fig. 13). To filter those samples with image-depth misalignment, we first run Canny edge detectors on both the image (converted to grayscale) and depth map to extract the boundaries. Next, we dilate the boundaries by 1 pixel and compute the percentage of intersection between image and depth boundaries. Then, we further dilate the boundaries by 3 pixels and compute the intersection between image and depth boundaries. Finally, we compute the ratio between the two intersections and remove samples whose ratio falls below a certain threshold.



Figure 13 Image-depth spatial misalignment in IRS dataset.

UnrealStereo4K. We remove the scene 00003 which consists of large erroneous regions. We also remove the scene 00004 which suffers from clipping issue. For scene 00008, we remove samples where the sea does not have depth values (images 9-13, 23-29, 80-82, 86-88, 96, 103-111, 126-136, 144-145, 148-154, 173-176, 178-179, 186-187, 191-192, 198-199) (Fig. 14).

GTA-SfM. We clip the maximum depth value at 1000.

Kenburns. We clip the depth value at 50,000 following GitHub issue ¹.

¹<https://github.com/sniklaus/3d-ken-burns/issues/40>



Figure 14 Erroneous samples in UnrealStereo4K dataset. The windows in scene 00003 are transparent; scene 00004 suffers from clipping issue; the sea in scene 00008 does not have depth values.

PointOdyssey. We remove the two scenes `animal2_s` and `dancingroom3_3rd` where the ground depth is incorrect².

TRELLIS. We remove all samples with suffix `stl`, `abc`, `STL`, `PLY`, `ply` as we noticed that these samples do not contain texture.

OmniObject3D. We clip the maximum depth value at 10.

²https://github.com/y-zheng18/point_odyssey/issues/6