

DepthART: Scaling Foundation Monocular Depth to Tiny Models

Feng Xue
University of Trento
Trento, Italy
feng.xue@unitn.it

Wu Chen
Beijing University of Posts
and Telecommunications
Beijing, China
chenw@bupt.edu.cn

Mingshuai Zhao
Guofeng Zhong
Beijing University of Posts
and Telecommunications
Beijing, China

Anlong Ming*
Beijing University of Posts
and Telecommunications
Beijing, China
mal@bupt.edu.cn

Haozhe Wang
The Hong Kong University
of Science and Technology
Hong Kong, China
jasper.whz@outlook.com

Dianqiao Lei
Tsinghua University
Beijing, China

Zhaowen Lin
Haiyang Zhang
Beijing University of Posts
and Telecommunications
Beijing, China

Nicu Sebe
University of Trento
Trento, Italy
niculae.sebe@unitn.it




Figure 1: DepthART brings Depth Anything-style capability to tiny models, achieving strong zero-shot generalization, reliable relative depth, stable metric scale, and high efficiency. DAnet and FastDepth are lightweight models trained on NYUD v2. Depth Anything v2-S uses a ViT-S encoder for relative depth, while Depth Anything v1-L is fine-tuned on NYUD v2. Affine-invariant depth is evaluated by per-image least-squares scale-and-shift fitting over valid pixels, whereas metric 3D point clouds are visualized *without any re-scaling* after fine-tuning. The plot compares standard-protocol NYUD v2 δ_1 with strict-FP32 model-only FPS. DepthART is profiled at 224^2 and 448^2 ; other methods use their standard runtime resolutions.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, Woodstock, NY

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXXX.XXXXXXX>

Abstract

Recent geometric foundation models (e.g., *Metric3D*, *Depth Anything* and *UniDepth*) have substantially improved monocular depth estimation (MDE) in both cross-scene generalization and metric-scale prediction, yet these gains have not translated to tiny models. We bridge this gap with **DepthART** (**Depth Anything Rethought for Tiny Models**), which is a compact MDE model for on-device deployment across diverse scenes. We first identify two capacity-driven bottlenecks in tiny models: (i) overfitting to dataset-specific distribution bias and (ii) unstable metric adaptation under camera shift, where full fine-tuning easily damages transferable geometry.

Accordingly, DepthART combines two simple but effective strategies: a bias-resistant data sampling scheme to reduce distribution bias under the same training budget, and a camera-conditioned fine-tuning protocol that freezes the distilled encoder and adjusts metric scale conditioned on intrinsics while better preserving cross-dataset generalization. Across datasets, DepthART consistently surpasses previous tiny baselines in both zero-shot generalization and metric accuracy (e.g., zero-shot $\delta_1=0.964$ for DepthART-S on NYUD v2), and in some cases approaches heavy models. We further provide a scalable model family, with DepthART-S reaching **347/245** FPS (strict FP32) on an RTX A6000 at $224^2/448^2$, **102** FPS (TF32) on a Orin NX 8GB, and over **15** FPS (FP32) on a Jetson Nano 4GB. Code is released on *GitHub*. Project page: xuefeng-cvr.github.io/DepthART.

CCS Concepts

• **Computing methodologies** → **3D imaging**; *Scene understanding*; *Reconstruction*.

Keywords

Foundation Monocular Depth Estimation, Tiny Model, Dataset Bias, Camera-conditioned Fine-tuning

ACM Reference Format:

Feng Xue, Wu Chen, Mingshuai Zhao, Guofeng Zhong, Anlong Ming, Haozhe Wang, Dianqiao Lei, Zhaowen Lin, Haiyang Zhang, and Nicu Sebe. 2026. DepthART: Scaling Foundation Monocular Depth to Tiny Models. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Monocular depth estimation (MDE) aims to recover 3D scene geometry from a single RGB image and serves as a core geometric primitive for downstream multimedia systems, such as image/video synthesis [16, 65, 67, 68], world models [3, 33], mobile robots [25, 61], and augmented reality [21]. As these applications move from controlled settings to real-world deployment, they increasingly require not only visually plausible depth, but also reliable generalization across diverse scenes and cameras.

To meet this demand, MDE has rapidly shifted toward foundation-model paradigms [13, 22, 29, 58, 66]. Early efforts improved generalization by training on broader mixtures (MiDaS [32]) and adopting stronger transformer backbones (DPT [31]). More recent work narrows the gap between relative-depth robustness and metric-scale prediction by scaling data and capacity and by introducing explicit camera modeling (Metric3D v1/v2 and UniDepth v1/v2 [13, 29, 30, 60]) or large-scale pseudo labels (Depth Anything v1/v2 [57, 58]). Some methods further expand the camera model beyond pinhole (UniK3D [28]) or predict richer 3D representations (MoGe v1/v2 [48, 49]). In parallel, generative priors offer a complementary route to robust dense geometry: diffusion-based approaches (e.g., Marigold [18] and GeoWizard [9]) often yield sharper structures and stronger robustness under distribution shift, and recent works accelerate inference via one-step transfer [12, 55] or flow matching [10]. Together, these advances substantially raise the ability ceiling of MDE.

However, these foundation capabilities have not become standard for tiny models that must run in real time on resource-limited devices. Rather than proposing a new camera-aware foundation depth architecture, we study a different regime: **how to preserve foundation-style transfer in tiny deployment-oriented backbones**. We find that naively shrinking foundation MDE fails for two regime-specific reasons: (i) **Overfitting to dataset-specific bias**: When trained on mixed data sources, tiny models are more likely to overfit dataset-specific bias (e.g., frequent patterns) rather than learn a depth prior that transfers across scenes. (ii) **Unstable metric adaptation under camera shift**: Tiny models often rely on implicitly encoded camera cues. As a result, full fine-tuning on a single metric dataset may improve in-domain accuracy, but it transfers poorly across cameras and more easily damages previously learned scene geometry.

To address this gap, we propose DepthART, a lightweight MDE model designed for real-time, on-device inference across diverse scenes. DepthART is built on two simple strategies aligned with the failure modes above. We first introduce a bias-resistant data sampling (BRDS) scheme. BRDS explicitly re-balances the long-tailed distribution of photographic and camera-induced cues that tend to dominate depth learning in ultra-light models, consistently improving cross-dataset zero-shot transfer under the same data budget. In Stage 1, we distill from Depth Anything v2 to transfer foundation-style relative depth priors to a tiny student. Then we design a camera-conditioned fine-tuning (CamFT) protocol. CamFT freezes the distilled encoder and inserts lightweight intrinsics-conditioned adapters plus a multi-query scale head to recover metric scale on the target domain, improving scale stability while substantially better preserving zero-shot geometry than full fine-tuning. Across extensive evaluations, DepthART substantially improves both scene generalization and metric-scale transfer over prior lightweight approaches, and in some settings approaches the performance of much larger models, as shown in the comparison of Fig. 1. We also provide multiple model sizes to match diverse deployment constraints, highlighting the practicality under real-time, resource-limited conditions. Our contributions are summarized as follows:

- We identify a practically important but underexplored regime: transferring foundation-style monocular depth generalization to tiny deployment-oriented models.
- We show that tiny MDE models are particularly vulnerable to dominant training distributions during distillation and to forgetting during metric fine-tuning, and address these issues with BRDS and encoder-frozen CamFT, respectively.
- With a low-latency TinyViM+DPT backbone (6M/11M/32M), DepthART substantially advances both cross-dataset relative-depth generalization and metric fine-tuning, offering a practical recipe for resource-limited deployment.

2 Related Work

Monocular Depth Estimation (MDE) [2, 7, 43, 52, 56, 69] was mainly studied in the single-domain supervised setting, focusing on architecture and loss design to recover scene structure from one dataset. Representative examples include DORN [8], VNL[59], and AdaBins [1], etc. These approaches achieved strong in-domain




Figure 2: Bias-resistant data sampling (BRDS) and distillation pipeline. A 44M multi-source corpus is encoded (VAE) and reduced (Incremental PCA) to enable density-aware sample selection, producing a 1.7M bias-resistant training pool. The scatter plots and density surfaces illustrate the effect of rebalancing in feature space and are simulated for visualization. DepthART (TinyViM-S/B/L + DPT) is then distilled from Depth Anything v2 Large using teacher depth supervision.

accuracy, but their generalization was often limited by the training data and camera settings. Recent progress is driven by foundation training. Methods such as MiDaS [32] and DPT [31] improve zero-shot relative depth through large-scale, mixed-data pretraining. Metric3D[13] and UniDepth[29] further target metric depth and camera variation. UniK3D[28] and MoGe [49] move toward more explicit 3D reasoning. Diffusion-based approaches (e.g., Lotus[12], Marigold[18], GeoWizard[9]) produce high-quality, detail-preserving depth (see introduction for more details). However, lightweight MDE methods, including FastDepth[51], DANet[40], GuideDepth[36], and LMDepth[23], are still trained and evaluated within a single domain, but with more efficient structure design. Therefore, they generally lag behind foundation models. DepthART is motivated by this gap: rather than proposing a new camera-aware foundation depth formulation, it aims to transfer foundation-level generalization and metric-scale robustness to compact models for on-device deployment.

Coreset Selection typically compresses training data by selecting a small *representative* subset that approximates full-data training, e.g., coverage-based k -center selection in feature space [38] or gradient/submodular matching methods such as CRAIG and Grad-Match [19, 26]. In monocular depth, distribution bias has been explicitly identified as a key obstacle to cross-domain generalization, e.g., DME [64] analyzes and mitigates depth-related bias. Inspired by both coreset selection and bias-focused depth studies, our sampling resembles coverage-based selection but targets debiasing rather than dataset compression, by suppressing over-dominant modes and promoting broad feature-space coverage for higher generalization.

3 Method

3.1 Overview of DepthART

DepthART targets tiny deployment-oriented MDE, with an emphasis on preserving cross-scene transfer and enabling stable metric calibration under strict model budgets. For efficiency, DepthART

adopts a compact Mamba-style encoder, *i.e.*, TinyViM [24], paired with a widely used DPT decoder [31], forming a lightweight trunk.

The training proceeds in two stages. In the first stage, we distill DepthART from a large foundation MDE model, *e.g.*, Depth Anything v2, to inherit robust cross-scene affine-invariant depth generalization. As tiny models are highly sensitive to dataset bias, we employ a bias-resistant data sampling scheme (*see Sec. 3.2*) to construct a bias-resistant subset from a large multi-source corpus for training. In the second stage, we freeze the distilled encoder and perform camera-conditioned fine-tuning (*see Sec. 3.3*) on a single metric dataset, *i.e.*, NYUD v2 or KITTI. To mitigate catastrophic forgetting, we insert a simple camera adapter to adjust the depth feature conditioned on camera intrinsics, and a multi-query scale head to regress the entire scale of images.

3.2 Bias-Resistant Relative Depth Pre-training

Dataset Distribution Bias Analysis. In practice, manually collected datasets follow implicit priors shaped by *geometry*, *viewpoint*, and *photometric* conditions. For example, **ImageNet** [6] is largely object-centric, **LSUN** [63] often captures indoor scenes from similar camera viewpoints, and **Google Landmark** [50] primarily focuses on landmarks, buildings, and natural landscapes, exhibiting a typical “tourist photography” bias and a very long-tailed category distribution. **BDD100K** [62], in contrast, exhibits strong ego-pose priors from vehicle-mounted cameras. However, comprehensive data coverage for MDE is more critical than tightly fitting frequent patterns of a single dataset. Therefore, it is essential to increase feature-space coverage and reduce the over-representation of high-density regions so that long-tail cases receive sufficient training exposure, especially for capacity-limited models.

Bias-Resistant Data Sampling Strategy. Given an RGB image set $\mathcal{D} = \{x_i\}_{i=1}^N$ from a dataset, our bias-resistant sampling applies density-aware stratified sub-sampling in an image-level feature space where similar images form clusters, down-weighting dominant modes to retain diversity (Fig. 2).

Specifically, for each image x_i , we extract a continuous representation $z_i = E(x_i)$ using a well-trained variational autoencoder (VAE) $E(\cdot)$ from Stable Diffusion v1.5 [35]. Unlike object-centric embeddings, e.g., DINO, VAE latents tend to emphasize holistic visual factors (e.g., layout, texture, illumination, and viewpoint) that are more aligned with bias in training depth. Secondly, we project z_i to a J -dimensional vector $\tilde{z}_i \in \mathbb{R}^J$ via Principal Component Analysis (PCA), denoted as $\tilde{z}_i = \text{PCA}(z_i)$, which yields a compact and smoothly structured embedding space suitable for large-scale sampling. Thirdly, to assign each sample to a grid cell, we perform per-dimension binning on $\tilde{z}_i$. For each dimension $j \in \{1, \dots, J\}$, we partition its value range into B bins with uniformly spaced boundaries $\{b_k^{(j)}\}_{k=0}^B$, where $b_0^{(j)} = \min_i \tilde{z}_i^{(j)}$ and $b_B^{(j)} = \max_i \tilde{z}_i^{(j)}$. We then compute J -D cell identifier, denoted as c_i , for an image x_i :

$$c_i = (c_i^{(1)}, \dots, c_i^{(J)}), c_i^{(j)} = \max \left\{ k \in [0, B-1] \mid b_k^{(j)} \leq \tilde{z}_i^{(j)} \right\}, \quad (1)$$

Images sharing the same c_i fall into the same grid cell and are treated as one stratum for subsequent density-aware sub-sampling. Subsequently, we perform a stratified sampling per cell. Let $\mathcal{D}_c = \{x_i | c_i = c\}$ denote the images that fall into cell c , the sub-sampled set, denoted as $\mathcal{S}_c \subset \mathcal{D}_c$, is formulated as:

$$\mathcal{S}_c \sim \mathcal{U}(\mathcal{D}_c, m_c), \quad m_c = \max(1, \lceil |\mathcal{D}_c| / \gamma \rceil), \quad (2)$$

where $\gamma > 1$ is the reduction factor, $\mathcal{U}(\mathcal{D}_c, m_c)$ denotes uniformly sub-sampling m_c images from $\mathcal{D}_c$. This rule enforces two properties: (i) **Coverage**: even sparse, underrepresented regions remain visible in $\mathcal{S}_c$ via $\max(1, \cdot)$ and (ii) **Bias reduction**: for dense, dominant regions, only a fraction of samples are retained, suppressing their influence. Finally, we aggregate all selected samples into a bias-resistant subset:

$$\mathcal{D}_{\text{BR}} = \bigcup_c \mathcal{S}_c. \quad (3)$$

In this way, $\mathcal{D}_{\text{BR}}$ preserves broad diversity across scenes and camera poses while suppressing the dominance of overrepresented modes in the original datasets.

Distillation from Foundation MDE Models. We distill DepthART on the offline sub-sampled set $\mathcal{D}_{\text{BR}}$ using *Depth Anything v2 Large* as the teacher. For each image in $\mathcal{D}_{\text{BR}}$ (across all training datasets), the teacher produces the pseudo depth and the student predicts the corresponding depth. We follow a standard MDE distillation [32]:

$$\mathcal{L}_{\text{Distill}} = \mathcal{L}_{\text{MSE}} + \alpha \mathcal{L}_{\text{Edge}}, \quad \alpha = 0.5, \quad (4)$$

where $\mathcal{L}_{\text{MSE}}$ measures the error between the teacher pseudo-depth and the student prediction after *per-image scale/shift alignment*, and $\mathcal{L}_{\text{Edge}}$ encourages consistent depth discontinuities via an L_1 penalty on spatial gradients. For higher robustness, $\mathcal{L}_{\text{MSE}}$ ignores the top-10% highest-error pixels in each image, preventing a few noisy pixels from dominating the distillation signal.

3.3 Camera-conditioned Metric Fine-tuning

Forgetting Effects in Tiny Models Fine-Tuning. Due to limited capacity, fully fine-tuning tiny models on a single metric dataset might easily overwrite the transferable depth cues learned via distillation, causing *catastrophic forgetting*.

We employ affine-invariant metric to measure the degradation of depth structure after fine-tuning. As shown in Fig. 3(a), full




Figure 3: Geometry and scale stability after fine-tuning on NYUD v2. (a) Full fine-tuning hurts zero-shot affine-invariant RMSE more than CamFT, indicating greater geometry damage. (b) On iBims-1, CamFT produces a tighter distribution of per-image scale-shift alignment parameters, indicating less variable per-image affine corrections than full fine-tuning.

fine-tuning on NYUD v2 improves in-domain accuracy (Affine-invariant RMSE: $0.27 \rightarrow 0.234$) but substantially degrades zero-shot performance on iBims-1 ($0.333 \rightarrow 0.421$), indicating that the depth structure of zero-shot scenes is damaged when all parameters are updated. Beyond structural degradation, fine-tuning also perturbs metric-scale. Fig. 3(b) shows increased variability in the per-image scale and shift alignment parameters on iBims-1 after full fine-tuning (std : $0.387/0.712$), reflecting unstable scale transfer.

Camera-conditioned Fine-Tuning. To adapt metric scale without overwriting the pre-trained depth prior, we perform a calibration-style fine-tuning that keeps the distilled encoder frozen (Fig. 4). Instead of fully fine-tuning the model, which easily overfits the target domain and damages learned geometry in tiny models, we add two trainable modules: (i) *camera adapters* that modulate encoder features using pyramid camera prompts, and (ii) a *multi-query scale head* that estimates a global scene scale from depth cues and camera prompt. Finally, the DPT decoder is trained to output relative depth $\hat{d}_{\text{rel}}$ re-scaled to $[0, 1]$, and the scale head predicts a positive scale factor, i.e., s , for metric calibration:

$$\hat{d}_{\text{metric}} = d_{\text{max}} (s \odot \hat{d}_{\text{rel}}). \quad (5)$$

where d_{max} is a constant maximum depth and $\odot$ denotes multiplication. We deliberately use a scale-only formulation in this work. This choice is not meant to imply that shift correction is universally unnecessary. Rather, in the tiny regime, we find it preferable to keep metric calibration simple and stable, so that adaptation does not interfere excessively with the transferred relative-depth prior. We next detail the *pyramid camera prompt*, *camera adapter*, and *multi-query scale head*.




Figure 4: Camera-conditioned fine-tuning. We freeze the distilled encoder and inject lightweight camera adapters conditioned on pyramid camera prompts. A multi-query scale head predicts a scale factor, transferring the relative depth to metric depth.

Pyramid Camera Prompt. To make metric fine-tuning camera-aware without sacrificing efficiency, we convert the camera intrinsics $\mathbf{K}=(f_x, f_y, c_x, c_y)$ into dense prompts for TinyViM. Following the ray-direction prompt construction in prior work [29], we compute a per-pixel camera-ray field from $\mathbf{K}$ and apply a Fourier encoding to obtain a prompt. Unlike single-resolution conditioning in fixed-resolution ViT pipelines, TinyViM progressively downsamples features ($1/4 \rightarrow 1/8 \rightarrow 1/16 \rightarrow 1/32$), where camera effects should be injected at matching resolutions. Accordingly, we build prompts $\{C^{(s)} \in \mathbb{R}^{C_s \times H/s \times W/s}\}$ at strides $s \in \{4, 8, 16, 32\}$, and inject them at the corresponding stages. In this way, high-resolution prompts guide local structure adaptation in shallow layers, while coarse prompts support global scale calibration in deep layers, making camera-conditioned fine-tuning more stable for tiny models.

Camera Adapter. Given the pyramid prompts, we inject them into each frozen encoder stage via a lightweight residual adapter. Let $\mathbf{F}^{(l)}$ be the frozen feature map at stage l . We compute an updated feature $\tilde{\mathbf{F}}^{(l)}$ as

$$\tilde{\mathbf{F}}^{(l)} = \mathbf{F}^{(l)} + \text{CA}(\mathbf{F}^{(l)}, \text{Proj}(C^{(s)})), \quad (6)$$

where CA is a 1-layer, 1-head cross-attention block with queries from $\mathbf{F}^{(l)}$ and keys/values from $\text{Proj}(C^{(s)})$. $\text{Proj}(\cdot)$ applies a 3×3 depthwise convolution followed by a 1×1 projection, without activation or normalization. We keep this adapter intentionally small to preserve real-time inference while enabling camera-conditioned metric calibration.

Multi-Query Scale Head. Estimating a single global scale from a tiny backbone is usually unstable, as the prediction may be affected by spurious textures or local outliers. We therefore design a multi-query scale head that aggregates multiple global cues via an implicit ensemble. Given the deepest feature $\tilde{\mathbf{F}}^{(4)}$ and its aligned camera prompt $C^{(32)}$, we form a context set $\mathbf{X} = [\tilde{\mathbf{F}}^{(4)}, C^{(32)}]$ by concatenation and apply a lightweight Transformer decoder with Q learnable queries. Specifically, with $\mathbf{Q} \in \mathbb{R}^{Q \times 128}$ and a 2-layer, 4-head decoder $\text{Dec}(\cdot)$, we compute the scale factor:

$$\mathbf{Z} = \text{Dec}(\mathbf{Q}, \mathbf{X}), \quad \mathbf{s} = \frac{1}{Q} \sum_{q=1}^Q \sigma(\text{MLP}(\mathbf{Z}_{q,:})), \quad (7)$$

where $\text{MLP}(\cdot)$ is a two-layer MLP ($128 \rightarrow 128 \rightarrow 1$) with ReLU, and $\sigma(\cdot)$ enforces a positive, bounded scale. Different queries attend

to complementary scale cues, e.g., scene layout/FoV-driven geometry/large planes. The averaging reduces their variance, yielding a more stable scale estimation.

Fine-tuning Phase. We perform CamFT on a single dataset (NYUD v2 or KITTI) with standard supervision. In this process, we minimize the SiLog loss on valid pixels with a multi-scale edge loss [32]:

$$\mathcal{L}_{\text{ft}} = \mathcal{L}_{\text{SiLog}}(\hat{\mathbf{d}}_{\text{metric}}, \mathbf{d}) + \lambda \mathcal{L}_{\text{Edge}}(\hat{\mathbf{d}}_{\text{metric}}, \mathbf{d}), \quad (8)$$

where $\mathbf{d}$ is the metric ground truth and λ weights the edge term. During CamFT, we freeze the *distilled encoder* and update the *camera adapters*, the *multi-query scale head*, and *DPT*, mitigating forgetting while calibrating metric scale.

4 Experiments

4.1 Dataset

To maximize coverage in relative-depth pre-training, we aggregate data from 11 sources: ApolloScapeExtra [15], BDD100K [62], Cityscapes [5], DIML-indoor [4], GoogleLandmark [50], Holopix50K [14], HRWSI [54], ImageNet21K [6], LSUN [63], Objects365 [39], and ReDWeb [53]. For the large-scale datasets (BDD100K, GoogleLandmark, ImageNet21K, LSUN, and Objects365), we do not use the full collections but use our BRDS to extract bias-resistant samples, resulting in a 1.7M image training set. After pre-training, we evaluate zero-shot affine-invariant depth on NYUD v2 (Eigen crop) [41], KITTI (Eigen crop) [45], ETH3D [37], DIODE [46], and DDAD [11]. For metric-scale fine-tuning, we train on a single metric dataset (NYUD v2 or KITTI with Eigen crop) and further report zero-shot metric transfer on iBims-1 [20], SUN RGB-D [42], DIODE, DDAD, and ETH3D.

4.2 Baselines

For *affine-invariant* depth, we compare with *MiDaS v3.1* [32], which provides a broad family of generalizable relative MDE models from lightweight to heavyweight variants, including a fastest model MiDaS-LeViT; *Metric3D v2* [13], as a representative generalizable metric MDE method; *Depth Anything v2* [58], as both a representative generalizable relative MDE method and our teacher model; and *AnyDepth* [34], which targets both lightweight deployment and cross-scene generalization. Since AnyDepth does not release full

Table 1: Zero-shot affine-invariant depth estimation. Better: AbsRel↓, δ_1 ↑. Best and second-best results are marked in bold and underlined. DepthART is reported at three scales: S/B/L. All latency is measured on an RTX A6000 with batch size 1 under strict FP32 execution with TF32 disabled, and should thus be interpreted as deployment-oriented throughput rather than a strict same-resolution speed benchmark. “*” marks in-domain result. “Params” denotes the full model parameter amount.

Method	Encoder	Params	Native Lat. A6000	NYUD v2		KITTI		ETH3D		DIODE-Full		DDAD	
				AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑
MiDaS v3.1 [32] <i>Comprehensive Models</i>	LeViT	51M	3.42 ms	0.090*	0.915*	0.127*	0.830*	0.129	0.875	0.240	0.698	0.127	0.836
	Swin2-T	42M	7.31 ms	0.075*	0.942*	0.103*	0.883*	0.119	0.886	0.229	0.718	0.117	0.848
	Swin2-L	213M	51.47 ms	0.055*	0.972*	0.078*	0.932*	0.086	0.949	0.212	0.747	0.094	0.899
	Beit-L	345M	156.56 ms	0.047*	0.980*	0.130*	0.843*	0.072	<u>0.961</u>	<u>0.198</u>	0.762	0.092	0.908
Metric3D v2 [13] <i>Large-scale Metric MDE</i>	ViT-S	37.5M	74.99 ms	0.056	0.965	<u>0.064</u>	<u>0.950</u>	<u>0.062</u>	0.955	0.247	<u>0.789</u>	–	–
	ViT-L	411.9M	452.54 ms	0.042	0.980	0.046	0.979	0.042	0.987	0.141	0.882	0.072	0.948
Depth Anything v2 [58] <i>Large-scale Relative MDE</i>	ViT-S	24.8M	13.06 ms	0.053	0.973	0.078	0.936	0.142	0.851	0.228	0.751	0.085	0.921
	ViT-B	97.5M	33.04 ms	0.049	0.976	0.078	0.939	0.137	0.858	0.236	0.756	<u>0.081</u>	<u>0.929</u>
	ViT-L	335.3M	102.46 ms	<u>0.045</u>	<u>0.979</u>	0.074	0.946	0.131	0.865	0.230	0.758	0.089	0.926
AnyDepth [34] <i>Real-world Deployable</i>	ViT-S	26.5M	–	0.082	0.932	0.102	0.883	0.084	0.935	0.247	0.714	–	–
	ViT-B	95.5M	–	0.072	0.950	0.097	0.901	0.080	0.945	0.236	0.727	–	–
	ViT-L	313.4M	–	0.060	0.968	0.086	0.926	0.096	0.954	0.226	0.736	–	–
DepthART (Ours) <i>Foundation Tiny MDE</i>	TinyViM-S	6.0M	4.09 ms	0.059	0.964	0.082	0.930	0.084	0.950	0.214	0.745	0.095	0.900
	TinyViM-B	11.4M	4.91 ms	0.057	0.969	0.088	0.929	0.092	0.954	0.214	0.746	0.094	0.906
	TinyViM-L	32.6M	8.07 ms	0.053	0.971	0.079	0.933	0.091	0.958	0.208	0.754	0.095	0.910

Table 2: In-domain (pink) and zero-shot (cyan) metric depth results without any re-scaling post-processing. Models are fine-tuned on NYUD v2 or KITTI with metric supervision and directly tested on unseen datasets. Best results are in bold. “–” denotes missing KITTI-fine-tuned checkpoints (e.g., FastDepth and DANet), so DDAD/ETH3D/DIODE results are not reported. Depth Anything v1-L is a heavy model.

Method	In Domain NYUD v2		Zero-shot				In Domain KITTI		Zero-shot				In Domain KITTI		Zero-shot			
	δ_1 ↑	RMSE↓	iBims-1 δ_1 ↑ RMSE↓	SUN RGB-D δ_1 ↑ RMSE↓	DIODE Indoor δ_1 ↑ RMSE↓		δ_1 ↑	RMSE↓	DDAD δ_1 ↑ RMSE↓	ETH3D Outdoor δ_1 ↑ RMSE↓			δ_1 ↑	RMSE↓	DIODE Outdoor δ_1 ↑ RMSE↓			
Depth Anything v1-L	0.984	0.206	0.639 0.965	0.647 0.585	0.291 2.504		0.982	1.90	0.787 8.214	0.164 6.507			0.288	8.994				
FastDepth	0.771	0.604	0.429 1.607	0.560 0.736	0.275 2.778		0.825 4.33	–	–	–	–	–	–	–	–	–	–	–
DANet	0.831	0.488	0.55 1.11	0.616 0.623	0.221 2.909		0.920 3.52	–	–	–	–	–	–	–	–	–	–	–
GuideDepth	0.840	0.478	0.560 1.014	0.603 0.657	0.271 2.766		0.882 4.10	0.197 21.31	0.171 9.539	0.158 14.25								
DepthART-S	0.924	0.340	0.713 0.693	0.809 0.379	0.292 1.687		0.943 3.12	0.811 7.140	0.228 6.518	0.216 9.862								
DepthART-B	0.942	0.307	0.732 0.672	0.831 0.342	0.287 2.018		0.951 2.948	0.847 6.475	0.227 6.831	0.231 9.474								
DepthART-L	0.946	0.295	0.747 0.646	0.847 0.326	0.298 2.012		0.956 2.784	0.867 6.340	0.258 5.794	0.249 8.823								

model code, we only report its parameter count and accuracy, not inference speed. We exclude diffusion-based methods because they are typically much slower, and self-supervised methods because they follow a different paradigm.

For *metric* depth, we compare with the lightweight models *FastDepth* [51], *DANet* [40], and *GuideDepth* [36]. We do not include *TuMDE* [44], *LightDepth* [17], or *METER* [27], because only partial code is publicly available and our reproduced results are substantially worse than those reported. We also include *Depth Anything v1-L* [57], which, although not lightweight, is the only publicly available Depth Anything model aligned with our setting. To keep the comparison fair, we only consider models fine-tuned on NYUD v2 or KITTI and tested on the remaining datasets. Accordingly, we do not treat large-scale camera-aware metric MDE models such as

UniDepth and Metric3D as direct baselines, since their training data, model capacity, and fine-tuning settings differ substantially from ours. Instead, our metric evaluation measures whether a model can retain useful metric transfer under strict deployment-oriented constraints. These large-scale models remain important related references, but they are not strict “apples-to-apples” comparisons in our setting.

4.3 Implementation Details

DepthART is trained on an A800, while all desktop latency is measured on an RTX A6000 with batch size 1 under strict FP32 execution with TF32 disabled. The main accuracy comparisons evaluate DepthART at 448×448 . For the accuracy–efficiency analysis in Table 6, we additionally report resolution-matched accuracy and




Figure 5: Qualitative comparison of affine-invariant depth and metric 3D reconstruction. Top-left: predictions on five zero-shot benchmarks (DDAD, DIODE, ETH3D, KITTI, and NYUD v2). Top-right: predictions on diverse self-collected scenes. Different representative baselines are used across the two groups to cover both lightweight and large models under limited space. DepthART-B more consistently preserves thin structures, object boundaries, and global scene layouts. Bottom: metric predictions after NYUD v2 fine-tuning are projected directly into 3D without post-hoc scale or shift alignment and evaluated on unseen indoor scenes. DepthART-S produces more coherent geometry around windows, doors, chairs, and furniture boundaries than DANet and Depth Anything v2-s (fine-tuned on Hypersim).

efficiency at 224×224 and at each method’s standard operating resolution. For BRDS, we use $J = 8$ and $B = 10$. Stage-1 distillation runs for 40 epochs with batch size 96 using Adam $(\beta_1, \beta_2) = (0.9, 0.999)$, learning rate $1e-4$, weight decay 0.01, and horizontal flips only. We apply source-specific reduction factors γ to the five large-scale datasets. CamFT runs for 40 epochs with batch size 64 and initial learning rate $5e-5$, with $Q = 8$, $d_{\max} = 10$ for NYUD v2, and $d_{\max} = 80$ for KITTI. We further apply mild intrinsics perturbation during CamFT by randomly rescaling images to $[0.9, 1.1]$ and cropping $[0.8, 1.0]$ of the resized area, which serves as regularization for small camera variations rather than severe intrinsics noise or missing metadata. RGB images use bilinear interpolation, while depth maps use nearest-neighbor interpolation.

4.4 Quantitative and Qualitative Comparison

Zero-Shot Affine-Invariant Depth Estimation. Table 1 shows zero-shot affine-invariant results on five benchmarks spanning diverse scenes. Across three lightweight scales (TinyViM-S/B/L), DepthART delivers consistently strong generalization on the same level of latency, outperforming early baselines (e.g., MiDaS) with orders-of-magnitude fewer parameters than ViT-L models. Notably, TinyViM-L achieves the best ETH3D δ_1 (0.958) among models below 50M parameters and improves AbsRel on DIODE-Full (0.208) over Depth Anything v2 ViT-L (0.230), showing that our BRDS helps to yield highly generalizable relative-depth cues. Although Metric3D v2 (ViT-L) attains the best accuracy, DepthART provides a far lighter, deployment-friendly alternative that better matches real-time, on-device constraints.

Table 3: Ablation of subset selection under three Stage-1 distillation budgets (0.5M/1.0M/1.7M). We compare BRDS with uniform random sampling and K-Means selection under the same distillation schedule. Bold marks the best result within each budget (ties allowed). Percentages show relative changes, with **red for degradation and **green** for improvement.**

Data Amount	Sampling Strategy	NYUD v2		KITTI	
		$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$
1.7M	Random	0.957 ^{-0.7%}	0.065 ^{+10.2%}	0.905 ^{-2.7%}	0.104 ^{+26.8%}
	K-Means	0.958 ^{-0.6%}	0.063 ^{+6.8%}	0.923 ^{-0.8%}	0.086 ^{+4.9%}
	Bias-Resistant	0.964	0.059	0.930	0.082
1.0M	Random	0.947 ^{-1.4%}	0.072 ^{+14.3%}	0.903 ^{-2.3%}	0.107 ^{+20.2%}
	K-Means	0.958 ^{-0.2%}	0.063 ^{-0.0%}	0.917 ^{-0.8%}	0.090 ^{+1.1%}
	Bias-Resistant	0.960	0.063	0.924	0.089
0.5M	Random	0.941 ^{-1.3%}	0.076 ^{+11.8%}	0.891 ^{-3.0%}	0.117 ^{+30.0%}
	K-Means	0.954 ^{+0.1%}	0.067 ^{-1.5%}	0.911 ^{-0.9%}	0.095 ^{+5.6%}
	Bias-Resistant	0.953	0.068	0.919	0.090

In-domain/Zero-shot Metric Depth Estimation. Table 2 shows both in-domain results (fine-tuned on NYUD v2 or KITTI) and cross-dataset transfer without further adaptation. Scaling DepthART from S to B/L generally improves both in-domain accuracy and cross-dataset transfer. After NYUD v2 fine-tuning, DepthART-L achieves the best results on iBims-1 (0.747/0.646) and SUN RGB-D (0.847/0.326). On DIODE-Indoor, DepthART-L obtains the best δ_1 (0.298), while DepthART-S achieves the lowest RMSE (1.687). After KITTI fine-tuning, DepthART-L achieves the best results on DDAD (0.867/6.340) and ETH3D-Outdoor (0.258/5.794), and the lowest RMSE on DIODE-Outdoor (8.823). Depth Anything v1-L retains the highest DIODE-Outdoor δ_1 (0.288). Overall, CamFT provides strong metric transfer across model scales while retaining competitive in-domain accuracy.

Qualitative results. Fig. 5 shows qualitative results for affine-invariant depth and metric 3D reconstruction. Across five zero-shot benchmarks and additional self-collected scenes, DepthART-B yields more coherent depth layouts and cleaner local structures than prior lightweight models, while remaining visually competitive with much larger baselines, i.e., the teacher, Depth Anything v2-L. It better preserves thin structures, object boundaries, and relative depth in challenging regions, where other methods often suffer from over-smoothing or local distortions. After metric fine-tuning, DepthART-S also produces more geometrically consistent 3D reconstructions on unseen indoor scenes, showing less distortion around the window and windowsill in the two left cases, and cleaner geometry for the wooden door and chair in the two right cases, compared with DANet and Depth Anything v2-S.

4.5 Additional Analysis

Ablations for Sampling Strategy. We ablate BRDS by distilling our model with the same sampling amount using Random, K-Means, or BRDS, under three budgets (0.5M/1.0M/1.7M). Table 3 reports zero-shot affine-invariant depth results on NYUD v2 and KITTI.

Table 4: CamFT Design ablations. Models are fine-tuned on NYUD v2 and evaluated in-domain (ID) on NYUD v2 and zero-shot (ZS) on iBims-1.

Frz	CamAdp	MQSH	NYUD v2 (ID)		iBims-1 (ZS)	
			$\delta_1 \uparrow$	RMSE $\downarrow$	$\delta_1 \uparrow$	RMSE $\downarrow$
×	×	×	0.928	0.328	0.572	0.844
✓	×	×	0.887	0.388	0.562	0.877
✓	✓	×	0.896	0.380	0.544	0.866
✓	×	✓	0.908	0.365	0.534	0.879
✓	✓	✓	0.924	0.340	0.713	0.693
×	✓	✓	0.930	0.333	0.613	0.779

Table 5: Ablations for adapter design. We study (i) how to fuse image features with camera prompts and (ii) where to insert the adapter. Cross-attention in the encoder yields the best in-domain accuracy on NYUD v2 and the strongest zero-shot transfer to iBims-1.

Adapter Design Choices	NYUD v2 (ID)		iBims-1 (ZS)	
	$\delta_1 \uparrow$	RMSE $\downarrow$	$\delta_1 \uparrow$	RMSE $\downarrow$
<i>Fusion method of camera prompt and feature</i>				
Additive	0.911 ^{-1.4%}	0.361 ^{+6.2%}	0.568 ^{-20.3%}	0.840 ^{+21.2%}
Concat + Proj	0.913 ^{-1.2%}	0.351 ^{+3.2%}	0.616 ^{-13.6%}	0.781 ^{+12.7%}
Cross-Attn	0.924	0.340	0.713	0.693
<i>Adapter insertion position</i>				
Decoder-side adapter	0.914 ^{-1.1%}	0.346 ^{+1.8%}	0.611 ^{-14.3%}	0.827 ^{+19.3%}
Encoder-side adapter	0.924	0.340	0.713	0.693

At 1.7M it achieves the best accuracy on both datasets (NYUD v2 0.964/0.059, KITTI 0.930/0.082), and under tighter budgets it remains more sample-efficient, notably improving KITTI over Random at 0.5M (δ_1 : 0.891 $\rightarrow$ 0.919, AbsRel: 0.117 $\rightarrow$ 0.090). These gains indicate that BRDS reduces dominant-mode overfitting and improves generalization with limited data. At 0.5M, BRDS is slightly behind K-Means on some metrics, as the budget is too small for the student to fully fit the more diverse, long-tail exposure emphasized by BRDS. With sufficient data (1.0M+), BRDS becomes consistently superior.

CamFT Design Ablations. Table 4 isolates three factors: freezing the distilled encoder (**Frz**), the camera adapter (**CamAdp**), and the Multi-Query Scale Head (**MQSH**). Without freezing, the model attains high in-domain accuracy on NYUD v2 (full fine-tuning: 0.928/0.328, or even 0.930/0.333 with CamAdp+MQSH) but transfers poorly to iBims-1 (0.572/0.844 and 0.613/0.779), showing severe cross-domain degradation. With the trunk frozen, removing either module causes a large drop on iBims-1 (w/o CamAdp: 0.534/0.879; w/o MQSH: 0.544/0.866), while using both achieves the best transfer (0.713/0.693) with comparable NYUD v2 results (0.924/0.340).

Adapter Design Choices. Table 5 evaluates the camera adapter design along two axes: fusion operator and insertion location. Replacing cross-attention with simple fusion between camera prompts and image features leads to clear degradation, especially under domain shift: on iBims-1, **Additive** drops from 0.713/0.693 to 0.568/0.840

Table 6: Accuracy–efficiency comparison on an RTX A6000 and on-device deployment results on a Jetson Orin NX. (a) reports efficiency under a controlled 224² setting and under the default operating settings of each method. (b) reports FP32 relative-depth inference on a Jetson Orin NX 8GB under MAXN mode. The best and second-best values are highlighted in bold and underlined, respectively.

(a) RTX A6000 48GB											(b) Orin NX 8GB					FP32 with TF32 enabled, MAXN				
Method	Params	Controlled Setting					Native Setting					Model	Input Res.	Peak Mem.	M-Lat.	E2E				
		All inputs are 224 ²					Swin2-T: 256 ² ; DA-v2: 518 ² ; DepthART: 448 ²													
		NYUD FLOPs	Mem.	M-Lat.	E2E		NYUD FLOPs	Mem.	M-Lat.	E2E										
		$\delta_1 \uparrow$	G $\downarrow$	MB $\downarrow$	ms $\downarrow$	ms $\downarrow$	$\delta_1 \uparrow$	G $\downarrow$	MB $\downarrow$	ms $\downarrow$	ms $\downarrow$									
MiDaS v3.1-LeViT	51.0	0.915 [*]	5.3755	213.31	3.42	10.10	N/A					Depth Anything v2-S	512 ²	124.71	126.40	127.15				
MiDaS v3.1-Swin2-T	42.0	N/A					0.942 [*]	27.58	242.65	7.31	13.66	DepthART-S	448 ²	59.29	29.48	40.40				
Depth Anything v2-S	24.8	N/A	7.735	125.82	4.83	10.05	0.973	41.304	228.60	13.06	21.70	DepthART-B	448 ²	<u>79.84</u>	<u>35.35</u>	<u>46.09</u>				
Depth Anything v2-B	97.5	N/A	29.874	515.85	8.52	14.08	<u>0.976</u>	159.507	694.55	33.04	40.77	DepthART-L	448 ²	162.74	61.22	67.10				
Depth Anything v2-L	335.3	N/A	109.778	1457.19	20.89	25.88	0.979	585.986	1770.14	102.46	145.71	Depth Anything v2-S	228 ²	113.38	18.94	21.65				
DepthART-S	6.0	0.953	1.749	46.36	2.88	3.72	0.964	6.996	91.15	4.09	6.98	DepthART-S	224 ²	50.68	9.80	12.43				
DepthART-B	11.4	<u>0.959</u>	<u>2.348</u>	<u>67.21</u>	<u>3.36</u>	<u>4.19</u>	0.969	<u>9.394</u>	<u>113.63</u>	<u>4.91</u>	<u>7.75</u>	DepthART-B	224 ²	<u>71.23</u>	<u>11.59</u>	<u>14.20</u>				
DepthART-L	32.6	0.967	6.178	152.29	5.23	6.08	0.971	24.711	210.39	8.07	10.93	DepthART-L	224 ²	152.72	18.02	20.63				

Notes. M-Lat. and E2E denote model-only and end-to-end latency, respectively. Mem. denotes peak CUDA-allocated memory. N/A indicates that the corresponding result is unavailable. OOM indicates that inference was terminated because of insufficient memory. * indicates that the model was trained or fine-tuned on the corresponding target dataset, rather than evaluated in a zero-shot manner. FPS is omitted because it is directly derived as 1000/M-Lat.




Figure 6: Model-only inference latency of DepthART on a Jetson Orin NX 8GB under different power modes. The three panels report relative-depth inference at 224² and 448², and metric-depth inference at 448², respectively. Each color represents one model scale (S, B, or L). TensorRT FP32 and TensorRT FP16 results are shown across the 10W, 15W, 20W, and MAXN modes, while PyTorch FP32 results are reported at MAXN only. FP32 execution enables TF32 acceleration. Lower latency is better.

and **Concat+Proj** to 0.616/0.781 (δ_1 /RMSE); even on NYUD v2, performance decreases to 0.911/0.361 (Additive) and 0.913/0.351 compared to 0.924/0.340. For insertion, placing the adapter on the decoder also hurts transfer, reducing iBims-1 from 0.713/0.693 to 0.611/0.827 (NYUD v2: 0.924/0.340 $\rightarrow$ 0.914/0.346). Therefore, we adopt an encoder-side cross-attention adapter in all experiments.

Efficiency Analysis. Table 6 reports strict-FP32 efficiency on an RTX A6000 and on-device deployment on a Orin NX 8GB. At 224², DepthART-S/B/L achieve 347/298/191 FPS with model-only latencies of 2.88/3.36/5.23 ms and end-to-end latencies of 3.72–6.08 ms. At 448², they retain 245/204/124 FPS with 4.09–8.07 ms model-only and 6.98–10.93 ms end-to-end latency. At their standard resolutions, DepthART-S/B/L are respectively 3.2 $\times$ /6.7 $\times$ /12.7 $\times$ faster than the corresponding Depth Anything V2 variants in model-only inference, while requiring substantially fewer FLOPs and less memory.

On the Orin NX, the relative-depth S/B/L models reach 102/86/55 FPS at 224² and 34/28/16 FPS at 448² in FP32 with TF32 enabled, with end-to-end latencies of 12.43–20.63 ms and 40.40–67.10 ms, respectively. Figure 6 further profiles relative- and metric-depth inference under the 10W, 15W, 20W, and MAXN modes. TensorRT consistently reduces latency across model scales and resolutions, with FP16 providing the highest throughput, while higher power modes yield further acceleration. Although incomplete Jetson Nano 4GB measurements are omitted from the table, DepthART-S still exceeds 15 FPS for 224² relative-depth inference in strict FP32 on it.

5 Conclusion

We present **DepthART**, a tiny monocular depth that transfers much of the generalization ability of foundation MDE models to deployment-oriented backbones, while supporting stable metric calibration through encoder-frozen camera-conditioned fine-tuning.

By combining bias-resistant sampling with camera-conditioned fine-tuning, DepthART improves zero-shot generalization and metric accuracy over prior lightweight baselines while maintaining high efficiency across a scalable model family.

Limitations and future work. (1) CamFT still follows the Depth Anything fine-tuning setting. Large-scale metric-depth training, e.g., in the spirit of UniDepth/Metric3D, is a key step to further strengthen metric-scale reliability. (2) CamFT assumes known camera intrinsics. We will explore intrinsics estimation, e.g., GeoCalib [47], to reduce this reliance and enable metric depth without meta-data.

References

- [1] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. 2021. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [2] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. 2023. ZoeDepth: Zero-Shot Transfer by Combining Relative and Metric Depth. arXiv:2302.12288 [cs.CV] <https://arxiv.org/abs/2302.12288>
- [3] Lucy Chai, Richard Tucker, Zhengqi Li, Phillip Isola, and Noah Snavely. 2023. Persistent Nature: A Generative Model of Unbounded 3D Worlds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [4] Jaehoon Cho, Dongbo Min, Youngjung Kim, and Kwanghoon Sohn. 2021. DIML/CVL RGB-D Dataset: 2M RGB-D Images of Natural Indoor and Outdoor Scenes. *arXiv preprint arXiv:2110.11590* (2021). arXiv:2110.11590
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [7] David Eigen and Rob Fergus. 2015. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *IEEE international conference on computer vision*.
- [8] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. 2018. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [9] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. 2024. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In *ECCV*.
- [10] Ming Gui, Johannes Schusterbauer, Ulrich Prestel, Pingchuan Ma, Dmytro Kotovenko, Olga Grebenkova, Stefan Andreas Baumann, Vincent Tao Hu, and Björn Ommer. 2025. DepthFM: Fast Generative Monocular Depth Estimation with Flow Matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [11] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 2020. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [12] Jing He, Haodong Li, Mingzhi Sheng, and Ying-Cong Chen. 2025. Lotus-2: Advancing Geometric Dense Prediction with Powerful Image Generative Model. arXiv:2512.01030 [cs.CV] <https://arxiv.org/abs/2512.01030>
- [13] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. 2024. 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* 46, 12 (2024), 10579–10596.
- [14] Yiwen Hua, Puneet Kohli, Pritish Uplavikar, Anand Ravi, Saravana Gunaseelan, Jason Orozco, and Edward Li. 2020. Holopix50k: A large-scale in-the-wild stereo image dataset. *arXiv preprint arXiv:2003.11172* (2020).
- [15] Xinyu Huang, Xinjing Cheng, Qichuan Geng, Binbin Cao, Dingfu Zhou, Peng Wang, Yuanqing Lin, and Ruigang Yang. 2018. The apollo-scape dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition workshop*.
- [16] Xueyang Kang, Zhengkang Xiang, Zezheng Zhang, and Kourosh Khoshelham. 2025. Look Beyond: Two-Stage Scene View Generation via Panorama and Video Diffusion. In *Proceedings of the ACM International Conference on Multimedia*. 9375–9384.
- [17] Fatemeh Baran Karimi, Amir Mehrpanah, and Reza Rawassizadeh. 2024. Light-Depth: A resource efficient depth estimation approach for dealing with ground truth sparsity via curriculum learning. *Robotics and Autonomous Systems* 181 (2024), 104784.
- [18] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. 2024. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [19] Krishnateja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, Abir De, and Rishabh Iyer. 2021. GRAD-MATCH: Gradient Matching based Data Subset Selection for Efficient Deep Model Training. In *Proceedings of the International Conference on Machine Learning*.
- [20] Tobias Koch, Lukas Liebel, Friedrich Fraundorfer, and Marco Korner. 2018. Evaluation of cnn-based single-image depth estimation methods. In *Proceedings of the European Conference on Computer Vision Workshops*.
- [21] Jaewook Lee, Filippo Aleotti, Diego Mazala, Guillermo Garcia-Hernando, Sara Vicente, Oliver James Johnston, Isabel Kraus-Liang, Jakub Powierza, Donghoon Shin, Jon E Froehlich, et al. 2025. Imaginatear: Ai-assisted in-situ authoring in augmented reality. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*.
- [22] Yihao Liu, Feng Xue, Anlong Ming, Mingshuai Zhao, Huadong Ma, and Nicu Sebe. 2024. SM4Depth: Seamless Monocular Metric Depth Estimation across Multiple Cameras and Scenes by One Model. In *Proceedings of the ACM International Conference on Multimedia*.
- [23] Jiahuan Long and Xin Zhou. 2025. LMDepth: Lightweight Mamba-based Monocular Depth Estimation for Real-World Deployment. *arXiv preprint arXiv:2505.00980* (2025). arXiv:2505.00980
- [24] Xiaowen Ma, Zhenliang Ni, and Xinghao Chen. 2025. Tinyvim: Frequency decoupling for tiny hybrid vision mamba. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- [25] Nicholas Merrill, Patrick Geneva, Michael Paton, Zixin Li, Yulin Yang, and Guoquan Huang. 2024. Fast and Robust Learned Single-View Depth-aided Monocular Visual-Inertial Initialization. *The International Journal of Robotics Research* 43, 2 (2024), 237–257.
- [26] Baharan Mirzasoleiman, Jeff Bilmes, and Jure Leskovec. 2020. Coresets for Data-efficient Training of Machine Learning Models. In *Proceedings of the International Conference on Machine Learning*.
- [27] Lorenzo Papa, Paolo Russo, and Irene Amerini. 2023. METER: A mobile vision transformer architecture for monocular depth estimation. *IEEE Transactions on Circuits and Systems for Video Technology* 33, 10 (2023), 5882–5893.
- [28] Luigi Piccinelli, Christos Sakaridis, Mattia Segu, Yung-Hsu Yang, Siyuan Li, Wim Abbeloos, and Luc Van Gool. 2025. UniK3D: Universal Camera Monocular 3D Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [29] Luigi Piccinelli, Christos Sakaridis, Yung-Hsu Yang, Mattia Segu, Siyuan Li, Wim Abbeloos, and Luc Van Gool. 2025. UniDepthV2: Universal Monocular Metric Depth Estimation Made Simpler. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025), 1–14.
- [30] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. 2024. UniDepth: Universal Monocular Metric Depth Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [31] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. 2021. Vision Transformers for Dense Prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 12179–12188. doi:10.1109/ICCV48922.2021.01196
- [32] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. 2022. Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-Shot Cross-Dataset Transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 3 (2022), 1623–1637. doi:10.1109/TPAMI.2020.3019967
- [33] Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, and Jun Gao. 2025. GEN3C: 3D-Informed World-Consistent Video Generation with Precise Camera Control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [34] Zeyu Ren, Zeyu Zhang, Wukai Li, Qingxiang Liu, and Hao Tang. 2026. AnyDepth: Depth Estimation Made Easy. *arXiv preprint arXiv:2601.02760* (2026). arXiv:2601.02760
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [36] Michael Rudolph, Youssef Dawoud, Ronja Gildenring, Lazaros Nalpanitidis, and Vasileios Belagiannis. 2022. Lightweight monocular depth estimation through guided decoding. In *Proceedings of the IEEE International Conference on Robotics and Automation*.
- [37] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. 2017. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [38] Ozan Sener and Silvio Savarese. 2018. Active Learning for Convolutional Neural Networks: A Core-Set Approach. In *International Conference on Learning Representations*.

- [39] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. 2019. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- [40] Fei Sheng, Feng Xue, Yicong Chang, Wenteng Liang, and Anlong Ming. 2022. Monocular depth distribution alignment with low computation. In *Proceedings of the IEEE International Conference on Robotics and Automation*.
- [41] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. 2012. Indoor segmentation and support inference from rgbd images. In *Proceedings of the European Conference on Computer Vision*.
- [42] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. 2015. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [43] Wen Su, Haifeng Zhang, Jia Li, Wenzhen Yang, and Zengfu Wang. 2019. Monocular depth estimation as regression of classification using piled residual networks. In *Proceedings of the ACM International Conference on Multimedia*. 2161–2169.
- [44] Xiaohan Tu, Cheng Xu, Siping Liu, Renfa Li, Guoqi Xie, Jing Huang, and Laurence Tianruo Yang. 2021. Efficient Monocular Depth Estimation for Edge Devices in Internet of Things. *IEEE Transactions on Industrial Informatics (TII)* 17, 4 (2021), 2821–2832.
- [45] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. 2017. Sparsity Invariant CNNs. In *Proceedings of the International Conference on 3D Vision*.
- [46] Igor Vasiljevic, Nick Kolkun, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z Dai, Andrea F Daniele, Mohammadreza Mostajabi, Steven Basart, Matthew R Walter, et al. 2019. Diode: A dense indoor and outdoor depth dataset. *arXiv preprint arXiv:1908.00463* (2019). arXiv:1908.00463
- [47] Alexander Veicht, Paul-Edouard Sarlin, Philipp Lindenberger, and Marc Pollefeys. 2024. GeoCalib: Single-image Calibration with Geometric Optimization. In *Proceedings of the European Conference on Computer Vision*.
- [48] Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. 2025. MoGe: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [49] Ruicheng Wang, Sicheng Xu, Yue Dong, Yu Deng, Jianfeng Xiang, Zelong Lv, Guangzhong Sun, Xin Tong, and Jiaolong Yang. 2025. MoGe-2: Accurate Monocular Geometry with Metric Scale and Sharp Details. arXiv:2507.02546 [cs.CV] <https://arxiv.org/abs/2507.02546>
- [50] Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. 2020. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [51] Diana Wofk, Fangchang Ma, Tien-Ju Yang, Sertac Karaman, and Vivienne Sze. 2019. Fastdepth: Fast monocular depth estimation on embedded systems. In *Proceedings of the IEEE International Conference on Robotics and Automation*.
- [52] Zizhang Wu, Zhuozheng Li, Zhi-Gang Fan, Yunzhe Wu, Jian Pu, and Xianzhi Li. 2023. V2Depth: Monocular depth estimation via feature-Level virtual-view simulation and refinement. In *Proceedings of the ACM International Conference on Multimedia*. 688–697.
- [53] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruiibo Li, and Zhenbo Luo. 2018. Monocular relative depth perception with web stereo data supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [54] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. 2020. Structure-Guided Ranking Loss for Single Image Depth Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [55] Guangkai Xu, Yongtao Ge, Mingyu Liu, Chengxiang Fan, Kangyang Xie, Zhiyue Zhao, Hao Chen, and Chunhua Shen. 2025. What Matters When Repurposing Diffusion Models for General Dense Perception Tasks?. In *International Conference on Learning Representations*.
- [56] Feng Xue, Junfeng Cao, Yu Zhou, Fei Sheng, Yankai Wang, and Anlong Ming. 2021. Boundary-induced and scene-aggregated network for monocular depth prediction. *Pattern Recognition* 115 (2021), 107901.
- [57] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. 2024. Depth Anything: Unleashing the Power of Large-Scale Unlabeled Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [58] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. 2024. Depth Anything V2. In *Advances in Neural Information Processing Systems*, Vol. 37. 21875–21911.
- [59] Wei Yin, Yifan Liu, and Chunhua Shen. 2021. Virtual normal: Enforcing geometric constraints for accurate and robust depth prediction. *TPAMI* 44, 10 (2021), 7282–7295.
- [60] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. 2023. Metric3D: Towards Zero-shot Metric 3D Prediction from A Single Image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9009–9019. doi:10.1109/ICCV51070.2023.00830
- [61] Naoki Yokoyama, Sehoon Ha, Dhruv Batra, Jiuguang Wang, and Bernadette Bucher. 2024. Vision-Language Frontier Maps for Zero-Shot Semantic Navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*.
- [62] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. 2020. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [63] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. 2015. LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop. *arXiv preprint arXiv:1506.03365* (2015).
- [64] Songsong Yu, Yifan Wang, Yunzhi Zhuge, Lijun Wang, and Huchuan Lu. 2024. DME: Unveiling the Bias for Better Generalized Monocular Depth Estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [65] Denis Zavadski, Johann-Friedrich Feiden, and Carsten Rother. 2024. ControlNet-XS: Rethinking the Control of Text-to-Image Diffusion Models as Feedback-Control Systems. In *Proceedings of the European Conference on Computer Vision*.
- [66] Longjian Zeng, Zunjie Zhu, Rongfeng Lu, Ming Lu, Bolun Zheng, Chenggang Yan, and Anke Xue. 2025. DepthDark: robust monocular depth estimation for low-light environments. In *Proceedings of the ACM International Conference on Multimedia*. 11239–11248.
- [67] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- [68] Wendong Zhang, Feng Gao, Bingbing Ni, Lingyu Duan, Yichao Yan, Jingwei Xu, and Xiaokang Yang. 2018. Depth structure preserving scene image generation. In *Proceedings of the ACM International Conference on Multimedia*. 727–736.
- [69] Kecheng Zheng, Zheng-Jun Zha, Yang Cao, Xuejin Chen, and Feng Wu. 2018. Lanet: Layout-aware dense network for monocular depth estimation. In *Proceedings of the ACM International Conference on Multimedia*. 1381–1388.