

UniDepth: Universal Monocular Metric Depth Estimation

Luigi Piccinelli¹ Yung-Hsu Yang¹ Christos Sakaridis¹
Mattia Segu¹ Siyuan Li¹ Luc Van Gool^{1,2} Fisher Yu¹

¹ETH Zürich ²INSAIT

Abstract

Accurate monocular metric depth estimation (MMDE) is crucial to solving downstream tasks in 3D perception and modeling. However, the remarkable accuracy of recent MMDE methods is confined to their training domains. These methods fail to generalize to unseen domains even in the presence of moderate domain gaps, which hinders their practical applicability. We propose a new model, UniDepth, capable of reconstructing metric 3D scenes from solely single images across domains. Departing from the existing MMDE methods, UniDepth directly predicts metric 3D points from the input image at inference time without any additional information, striving for a universal and flexible MMDE solution. In particular, UniDepth implements a self-promptable camera module predicting dense camera representation to condition depth features. Our model exploits a pseudo-spherical output representation, which disentangles camera and depth representations. In addition, we propose a geometric invariance loss that promotes the invariance of camera-prompted depth features. Thorough evaluations on ten datasets in a zero-shot regime consistently demonstrate the superior performance of UniDepth, even when compared with methods directly trained on the testing domains. Code and models are available at: github.com/lpiccinelli-eth/unidepth.

1. Introduction

The precise pixel-wise depth estimation is crucial to understanding the geometric scene structure, with applications in 3D modeling [9], robotics [10, 52], and autonomous vehicles [31, 41]. However, delivering reliable metric scaled depth outputs is necessary to perform 3D reconstruction effectively, thus motivating the challenging and inherently ill-posed task of Monocular Metric Depth Estimation (MMDE).

While existing MMDE methods [2, 13, 15, 33, 34, 36, 50] have demonstrated remarkable accuracy across different benchmarks, they require training and testing on datasets with similar camera intrinsics and scene scales. Moreover, the training datasets typically have a limited size and contain little diversity in scenes and cameras. These characteris-

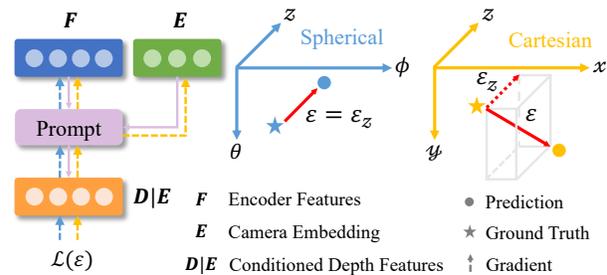


Figure 1. We introduce UniDepth, a novel approach that directly predicts 3D points in a scene with only one image as input. UniDepth incorporates a camera self-prompting mechanism and leverages a pseudo-spherical 3D output space defined by azimuth and elevation angles, and depth (θ , ϕ , z). This design effectively separates camera and depth optimization by avoiding gradient flowing to the camera module due to depth-related error (ϵ_z).

tics result in poor generalization to real-world inference scenarios [42], where images are captured in uncontrolled, arbitrarily structured environments and cameras with arbitrary intrinsics.

Only a few methods [20, 48] have addressed the challenging task of generalizable MMDE. However, these methods assume controlled setups at test time, including camera intrinsics. While this assumption simplifies the task, it has two notable drawbacks. Firstly, it does not address the full application spectrum, *e.g.* in-the-wild video processing and crowd-sourced image analysis. Secondly, the inherent camera parameter noise is directly injected into the model, leading to large inaccuracies in the high-noise case.

In this work, we address the more demanding task of generalizable MMDE *without* any reliance on additional external information, such as camera parameters, thus defining the universal MMDE task. Our approach, named UniDepth, is the first that attempts to solve this challenging task without restrictions on scene composition and setup and distinguishes itself through its general and adaptable nature. Unlike existing methods, UniDepth delivers metric 3D predictions for any scene *solely* from a single image, waiving the need for extra information about scene or camera. Further-

more, UniDepth flexibly allows for the incorporation of additional camera information at test time.

Our design introduces a camera module that outputs a non-parametric, *i.e.* dense camera representation, serving as the prompt to the depth module. However, relying only on this single additional module clearly results in challenges related to training stability and scale ambiguity. We propose an effective pseudo-spherical representation of the output space to disentangle the camera and depth dimensions of this space. This representation employs azimuth and elevation angle components for the camera and a radial component for the depth, forming a perfect orthogonal space between the camera plane and the depth axis. Moreover, the camera components are embedded through Laplace spherical harmonic encoding. Figure 1 depicts our camera self-prompting mechanism and the output space. Additionally, we introduce a geometric invariance loss to enhance the robustness of depth estimation. The underlying idea is that the camera-conditioned depth features from two views of the same image should exhibit reciprocal consistency. In particular, we sample two geometric augmentations, creating a pair of different views for each training image, thus simulating different apparent cameras for the original scene.

Our overall contribution is the first universal MMDE method, UniDepth, that predicts a point in metric 3D space for each pixel without *any* input other than a single image. In particular, first, we design a promptable camera module, an architectural component that learns a dense camera representation and allows for non-parametric camera conditioning. Second, we propose a pseudo-spherical representation of the output space, thus solving the intertwined nature of camera and depth prediction. In addition, we introduce a geometric invariance loss to disentangle the camera information from the underlying 3D geometry of the scene. Moreover, we extensively test UniDepth and re-evaluate seven MMDE State-of-the-Art (SotA) methods on ten different datasets in a fair and comparable zero-shot setup to lay the ground for the generalized MMDE task. Owing to its design, UniDepth consistently sets the new state of the art even compared with non-zero-shot methods, ranking first in the competitive official KITTI Depth Prediction Benchmark.

2. Related Work

Metric and Scale-agnostic Depth Estimation. It is crucial to distinguish Monocular Metric Depth Estimation (MMDE) from scale-agnostic, namely up-to-a-scale, monocular depth estimation. MMDE SotA approaches typically confine training and testing to the same domain. However, challenges arise, such as overfitting to the training scenario leading to considerable performance drops in the presence of minor domain gaps, often overlooked in benchmarks like NYU-Depthv2 [28] (NYU) and KITTI [17]. On the other hand, scale-agnostic depth methods, including MiDaS [35], Om-

niData [12], and LeReS [47], show robust generalization by training on extensive datasets. Their limitation lies in the absence of a metric output, hindering practical usage in downstream applications.

Monocular Metric Depth Estimation. The introduction of end-to-end trainable neural networks in MMDE, pioneered by [13], marked a significant milestone, also introducing the optimization process through the Scale-Invariant log loss (SI_{\log}). Subsequent developments witnessed the emergence of advanced networks, ranging from convolution-based architectures [15, 23, 25, 33] to transformer-based approaches [2, 34, 46, 50]. Despite impressive achievements on established benchmarks, MMDE models face challenges in zero-shot scenarios, revealing the need for robust generalization against domain shifts in appearance and geometry.

General Monocular Metric Depth Estimation. Recent efforts focus on developing MMDE models [3, 20, 48] for general depth prediction across diverse domains. These models often leverage camera awareness, either by directly incorporating external camera parameters into computations [14, 20] or by normalizing the shape or output depth based on intrinsic properties, as seen in [1, 24, 48].

However, these generalizable MMDE methods often adopt specific strategies to enhance performance, *e.g.* geometric pretraining [3] or dataset-specific prior like reshaping [48]. In addition, these methods assume access to noiseless camera intrinsics both at training and test time, also limiting their applicability to pinhole camera models. Additionally, SotA methods depend on a predefined backprojection operation, blurring the distinction between learning depth and the 3D scene. In contrast, our approach aims to overcome these limitations, presenting a more demanding perspective, *e.g.* universal MMDE. Universal MMDE involves directly predicting the 3D scene from the input image *without any* additional information other than the latter. Notably, we do not require any additional prior information at test time, such as access to camera information.

3. UniDepth

MMDE SotA methods typically assume access to the camera intrinsics, thus blurring the line between pure depth estimation and actual 3D estimation. In contrast, UniDepth aims to create a universal MMDE model deployable in diverse scenarios without relying on any other external information, such as camera intrinsic, thus leading to 3D space estimation by design. However, attempting to directly predict 3D points from a single image without a proper internal representation neglects geometric prior knowledge, *i.e.* perspective geometry, burdening the learning process with re-learning laws of perspective projection from data.

Sec. 3.1 introduces a pseudo-spherical representation of the output space to inherently disentangle camera rays' angles from depth. In addition, our preliminary studies indi-

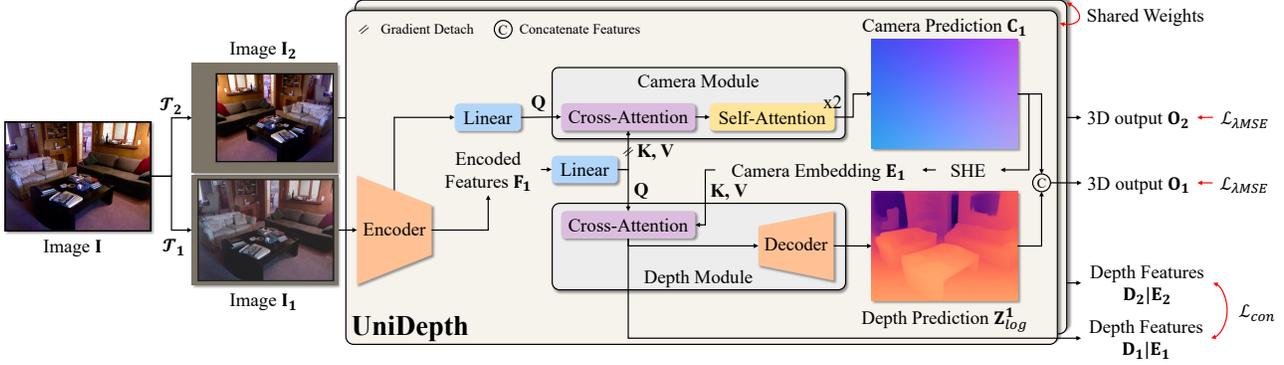


Figure 2. **Model Architecture.** UniDepth utilizes solely the input image to generate the 3D output (\mathbf{O}). It bootstraps dense camera prediction (\mathbf{C}) from the Camera Module, injecting prior knowledge on scene scale into the Depth Module via a cross-attention layer. The camera representation corresponds to azimuth and elevation angles. The geometric invariance loss (\mathcal{L}_{con}) enforces consistency between depth features tensors conditioned on the camera from different geometric augmentations ($\mathcal{T}_1, \mathcal{T}_2$). Stop-gradient is applied to the encoded feature (\mathbf{F}) flowing to the Camera Module to prevent the camera gradient from dominating the depth gradient in the encoder. The depth output (\mathbf{Z}_{log}) is obtained through three self-attention blocks interleaved with learnable 2x upsampling. The final output is the concatenation of the camera and depth tensors ($\mathbf{C}||\mathbf{Z}_{\text{log}}$), creating two independent optimization spaces for $\mathcal{L}_{\lambda\text{MSE}}$.

cate that depth prediction clearly benefits from prior information on the acquisition sensor, leading to the introduction of a self-prompting camera operation in Sec. 3.2. Further disentanglement at the level of internal depth features is achieved through a geometric invariance loss, outlined in Sec. 3.3. This loss ensures depth features remain invariant when conditioned on the bootstrapped camera predictions, promoting robust camera-aware depth predictions. The overall architecture and the resulting optimization induced by the combination of design choices are detailed in Sec. 3.4.

3.1. 3D Representation

The general purpose nature of our MMDE method requires inferring both depth and camera intrinsics to make 3D predictions based only on imagery observations. We design the 3D output space presenting a natural disentanglement of the two sub-tasks, namely depth estimation and camera calibration. In particular, we exploit the pseudo-spherical representation where the basis is defined by azimuth, elevation, and log-depth, *i.e.* $(\theta, \phi, z_{\text{log}})$, in contrast to the Cartesian representation (x, y, z) . The strength of the proposed pseudo-spherical representation lies in the decoupling of camera (θ, ϕ) and depth (z_{log}) components, ensuring their orthogonality by design, in contrast to the entanglement present in Cartesian representation.

It is worth highlighting that in this output space, the non-parametric dense representation of the camera is mathematically represented as a tensor $\mathbf{C} \in \mathbb{R}^{H \times W \times 2}$, where H and W are the height and width of the input image and the last dimension corresponds to azimuth and elevation values. While in the typical Cartesian space, the backprojection involves the multiplication of homogeneous camera rays and depth, the backprojection operation in the proposed representation space accounts for the concatenation of camera and depth representations. The pencil of rays are defined

as $(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) = \mathbf{K}^{-1}[\mathbf{u}, \mathbf{v}, \mathbf{1}]^T$, where \mathbf{K} is the calibration matrix, \mathbf{u} and \mathbf{v} are pixel positions in pixel coordinates, and $\mathbf{1}$ is a vector of ones. Therefore, the homogeneous camera rays $(\mathbf{r}_x, \mathbf{r}_y)$ correspond to $(\frac{\mathbf{r}_1}{\mathbf{r}_3}, \frac{\mathbf{r}_2}{\mathbf{r}_3})$.

Moreover, the angular dense representation can be embedded via the Laplace Spherical Harmonic Encoding (SHE). The camera embedding tensor is defined as $\mathbf{E} = \text{SHE}(\mathbf{C})$, $\mathbf{E} \in \mathbb{R}^{H \times W \times d}$, where d is the number of harmonics chosen. $\text{SHE}(\cdot)$ computes the set of spherical harmonics, *i.e.*, $\{\mathcal{Y}\}_{l,m}$ with degree l and order m , and concatenating along the channel dimension, with \mathbf{Y}_m^l as

$$Y_m^l(\theta, \phi) = \alpha_m^l \mathcal{P}_m^l(\cos \theta) e^{im\phi}, \quad (1)$$

where \mathcal{P}_m^l is the associated Legendre polynomial of degree l and order m , and α_m^l is a normalizing constant. In particular, the spherical harmonics on the unit sphere form an orthogonal basis of the spherical manifold and preserve inner products. The total number of harmonics utilized is 81, resulting from capping the degree l to 8. SHE is utilized as a mathematic sounder choice compared to, *e.g.* the Fourier Transform, to produce the camera embeddings.

3.2. Self-Promptable Camera

The camera module plays a crucial role in the final 3D predictions since its angular dense output accounts for two dimensions of the output space, namely azimuth and elevation. Most importantly, these embeddings prompt the depth module to ensure a bootstrapped prior knowledge of the input scene’s global depth scale. The prompting is fundamental to avoid mode collapse in the scene scale and to alleviate the depth module from the burden of predicting depth from scratch as the scale is already modeled by camera output.

Nonetheless, the internal representation of the camera module is based on a pinhole parameterization, namely via focal length (f_x, f_y) and principal point (c_x, c_y) . The four

tokens conceptually corresponding to the intrinsics are then projected to scalar values, *i.e.*, $\Delta f_x, \Delta f_y, \Delta c_x, \Delta c_y$. However, they do not directly represent the camera parameters, but the multiplicative residuals to a pinhole camera initialization, namely $\frac{H}{2}$ for y-components and $\frac{W}{2}$ for x-components, leading to $f_x = \frac{\Delta f_x W}{2}$, $f_y = \frac{\Delta f_y H}{2}$, $c_x = \frac{\Delta c_x W}{2}$, $c_y = \frac{\Delta c_y H}{2}$, leading to invariance towards input image sizes.

Subsequently, a backprojection operation based on the intrinsic parameters is applied to every pixel coordinate to produce the corresponding rays. The rays are normalized and thus represent vectors on a unit sphere. The critical step involves extracting azimuth and elevation from the backprojected rays, effectively creating a “dense” angular camera representation. This dense representation undergoes SHE to produce the embeddings \mathbf{E} . The embedded representations are then seamlessly passed to the depth module as a prompt, where they play a vital role as a conditioning factor. The conditioning is enforced via a cross-attention layer between the initialized feature of Depth Module $\mathbf{D} \in \mathbb{R}^{h \times w \times C}$ and the camera embeddings \mathbf{E} where $(h, w) = (H/16, W/16)$. The camera-prompted depth features $\mathbf{D}|\mathbf{E} \in \mathbb{R}^{h \times w \times C}$ are defined as

$$\mathbf{D}|\mathbf{E} = \text{MLP}(\text{CA}(\mathbf{D}, \mathbf{E})), \quad (2)$$

where CA is a cross-attention block and MLP is a Multi-Layer Perceptron with one 4C-channel hidden layer.

Figure 3 illustrates one of the main benefits of our camera module. In particular, in high-noise intrinsics or camera-agnostic scenarios, UniDepth can bootstrap the camera prediction, thus displaying total noise insensitivity. However, we can substitute the camera module output to improve 3D reconstruction peak performance if any external dense camera representation is provided. This adaptability enhances the model’s versatility, allowing it to operate seamlessly in diverse setups. Moreover, Figure 3 suggests that training with noisy self-prompts enhances the robustness of UniDepth to noisier external intrinsics if given at test time.

3.3. Geometric Invariance Loss

The spatial locations from the same scene captured by different cameras should correspond when the depth module is conditioned on the specific camera. To this end, we propose a geometric invariance loss to enforce the consistency of camera-prompted depth features of the same scene from different acquisition sensors. In particular, consistency is enforced on features extracted from identical 3D locations.

For each image, we perform N distinct geometrical augmentations, denoted as $\{\mathcal{T}_i\}_{i=1}^N$, with $N = 2$ in our experiments. This operation involves involves sampling a rescaling factor $r \sim 2^{\mathcal{U}[-1,1]}$ and a relative translation on the x -axis $t \sim \mathcal{U}[-0.1,0.1]$, then cropping it to the network’s input shape. This is analogous to sampling a pair of images from the same scene and extrinsic parameters but captured by different cameras. Let \mathbf{C}_i and $\mathbf{D}_i|\mathbf{E}_i$ describe the predicted

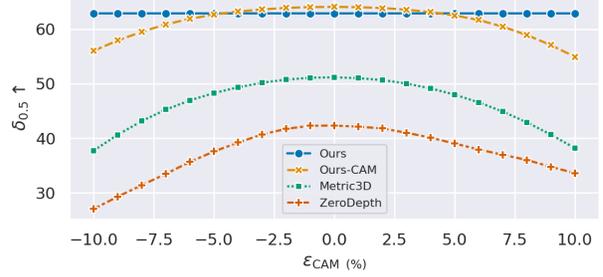


Figure 3. **Impact of noise in camera intrinsics.** The amount of relative distortion ($\epsilon_{CAM}(\%)$) of the intrinsics is shown on the x-axis, while $\delta_{0.5}$ performance on OOD test sets on the y-axis. Relying on external input inherently leads to being subject to its noise. UniDepth functions in dual regimes, with and without external intrinsic. In situations of unknown intrinsics or high noise, UniDepth exhibits total robustness by bootstrapping camera prediction (*Ours*). In contrast, with low-noise intrinsics, we leverage it for enhanced peak performance (*Ours-CAM*).

camera representation and camera-prompted depth features, respectively, corresponding to augmentation \mathcal{T}_i . It is evident that the camera representations differ when two diverse geometric augmentations are applied, *i.e.*, $\mathbf{C}_i \neq \mathbf{C}_j$ if $\mathcal{T}_i \neq \mathcal{T}_j$. Therefore, the geometric invariance loss can be expressed as

$$\mathcal{L}_{\text{con}}(\mathbf{D}_1|\mathbf{E}_1, \mathbf{D}_2|\mathbf{E}_2) = \left\| \mathcal{T}_2 \circ \mathcal{T}_1^{-1} \circ (\mathbf{D}_1|\mathbf{E}_1) - \text{sg}(\mathbf{D}_2|\mathbf{E}_2) \right\|_1, \quad (3)$$

where $\mathbf{D}_i|\mathbf{E}_i$ represents the depth feature after being conditioned by camera prompt \mathbf{E}_i , as outlined in Sec. 3.2, and $\text{sg}(\cdot)$ corresponds to the stop-gradient detach operation needed to exploit $\mathbf{D}_2|\mathbf{E}_2$ as pseudo ground-truth (GT). The bi-directional loss can be computed as: $\frac{1}{2}(\mathcal{L}_{\text{con}}(\mathbf{D}_1|\mathbf{E}_1, \mathbf{D}_2|\mathbf{E}_2) + \mathcal{L}_{\text{con}}(\mathbf{D}_2|\mathbf{E}_2, \mathbf{D}_1|\mathbf{E}_1))$. It is necessary to apply the geometric invariance loss *after* the features are conditioned on the viewing information, *i.e.*, camera. Otherwise, the loss would enforce consistency across features that inherently carry distinct camera information.

3.4. Network Design

Architecture. Our network, described in Fig. 2, comprises an Encoder Backbone, a Camera Module, and a Depth Module. The encoder can be either convolutional or ViT-based [11], producing features at different “scales”, *i.e.* $\mathbf{F} \in \mathbb{R}^{h \times w \times C \times B}$, where $(h, w) = (\frac{H}{16}, \frac{W}{16})$ and $B = 4$.

The Camera Module parameters are initialized class tokens for ViT-style or pooled feature maps for convolutional-style backbones. The encoded features from the Encoder Backbone are passed to the Camera Module as a stack of detached tokens, the encoder class tokens are utilized as camera parameters initialization. The features are processed to obtain the final dense representation \mathbf{C} as detailed in Sec. 3.2, and further embedded to \mathbf{E} via $\text{SHE}(\cdot)$ outlined in Sec. 3.1. Note that the stop-gradient operation is necessary because

of the low variety of effective cameras compared to the image diversity. In fact, the Camera Module component easily overfits and clearly dominates the overall backbone gradient.

The Depth Module is fed with the encoder features to condition the initial latent features $\mathbf{L} \in \mathbb{R}^{h \times w \times C}$ via one cross-attention layer to obtain the initial depth features, \mathbf{D} . The latent feature tensor \mathbf{L} is obtained as the average of the features \mathbf{F} along the B dimension. Furthermore, the depth features are conditioned on the camera prompts \mathbf{E} to obtain $\mathbf{D}|\mathbf{E}$ as described in Sec. 3.2. The camera-prompted depth features are further processed via self-attention layers where the positional encoding utilized is \mathbf{E} and upsampled to produce a multi-scale output. The log-depth prediction $\mathbf{Z}_{\log} \in \mathbb{R}^{H \times W \times 1}$ corresponds to the mean of the interpolated intermediate representations. The final 3D output $\mathbf{O} \in \mathbb{R}^{H \times W \times 3}$ is the concatenation of predicted rays and depth, $\mathbf{O} = \mathbf{C} \parallel \mathbf{Z}$, with \mathbf{Z} as element-wise exponentiation of \mathbf{Z}_{\log} .

Optimization. The optimization process is guided by a reformulation of the Mean Squared Error (MSE) loss in the final 3D output space (θ, ϕ, z_{\log}) from Sec. 3.1 as:

$$\mathcal{L}_{\lambda\text{MSE}}(\varepsilon) = \|\mathbb{V}[\varepsilon]\|_1 + \boldsymbol{\lambda}^T (\mathbb{E}[\varepsilon] \odot \mathbb{E}[\varepsilon]), \quad (4)$$

where $\varepsilon = \hat{\mathbf{o}} - \mathbf{o}^* \in \mathbb{R}^3$, $\hat{\mathbf{o}} = (\hat{\theta}, \hat{\phi}, \hat{z}_{\log})$ is the predicted 3D output, $\mathbf{o}^* = (\theta^*, \phi^*, z_{\log}^*)$ is the GT 3D value, and $\boldsymbol{\lambda} = (\lambda_{\theta}, \lambda_{\phi}, \lambda_z) \in \mathbb{R}^3$ is a vector of weights for each dimension of the output. $\mathbb{V}[\varepsilon]$ and $\mathbb{E}[\varepsilon]$ are computed as the vectors of empirical variances and means for each of the three output dimensions over all pixels, *i.e.* $\{\varepsilon^{(i)}\}_{i=1}^N$. Note that if $\lambda_d = 1$ for a dimension d , the loss represents the standard MSE loss for that dimension. If $\lambda_d < 1$, a scale-invariant loss term is added to that dimension if it is expressed in log space, *e.g.* for the depth dimension z_{\log} and a shift-invariant loss term is added if that output is expressed in linear space. In particular, if only the last output dimension is considered, *i.e.*, the one corresponding to depth, and $\lambda_z = 0.15$ is utilized, the corresponding loss is the standard SI_{\log} . In our experiments, we set $\lambda_{\theta} = \lambda_{\phi} = 1$ and $\lambda_z = 0.15$. Therefore, the final optimization loss is defined as

$$\mathcal{L} = \mathcal{L}_{\lambda\text{MSE}} + \alpha \mathcal{L}_{\text{con}}, \quad \text{with } \alpha = 0.1. \quad (5)$$

The loss defined here serves as a motivation for the designed output representation. Specifically, employing a Cartesian representation and applying the loss directly to the output space would result in backpropagation through (x, y) , and z_{\log} errors. However, x and y components are derived as $r_x \cdot z$ and $r_y \cdot z$ as detailed in Sec. 3.1. Consequently, the gradients of camera components, expressed by (r_x, r_y) , and of depth become intertwined, leading to suboptimal optimization as discussed in Sec. 4.3.

4. Experiments

4.1. Experimental Setup

In-domain training datasets. The training dataset utilized is the ensemble of Argoverse2 [43], Waymo [39], Driv-

ingStereo [45], Cityscapes [6], BDD100K [49], Mapillary-PSD [1], A2D2 [18], ScanNet [7], and Taskonomy [51]. The resulting dataset amounts roughly to 3M real-world images with different cameras and domains, compared to, *e.g.* Metric3D [48] and ZeroDepth [20] which exploit 8M and 17M training images, respectively.

Zero-shot testing datasets. We evaluate the generalizability of the compared models by testing them on ten datasets not seen during training. More precisely, each method is tested on validation splits from SUN-RGBD [38] without NYU split, Diode Indoor [40], IBims-1 [22], VOID [44] HAMMER [21], ETH-3D [37], nuScenes [4], and DDAD [19] with split proposed in [34] and evaluated with official masks. Also, UniDepth and the models from [20, 48] are zero-shot-tested on NYU-Depth V2 [28] and KITTI [17]. In particular, KITTI testing is performed on the corrected Eigen-split test set [13] with the Garg evaluation mask [16], while NYU testing uses the evaluation mask from [24].

Evaluation Details. All methods have been re-evaluated with a fair and consistent pipeline. In particular, we do not exploit any test-time augmentations. We use training image shapes for zero-shot testing and evaluate on the same validation splits and masks. Unfortunately, ZeroDepth lacks full code reproducibility, thus we report results from the original paper only, and for visualization, we utilize their provided code and weights. When methods do not report the configuration for a specific test dataset, we use the settings of NYU and KITTI for indoor and outdoor testing, respectively. We utilize common depth estimation evaluation metrics: root mean square error (RMS) and its log variant (RMS_{\log}), absolute mean relative error (A.Rel), the percentage of inlier pixels (δ_i) with threshold 1.25^i , scale-invariant error in log-scale (SI_{\log}): $100\sqrt{\text{Var}(\varepsilon_{\log})}$. In addition, we report point-cloud-based metrics proposed in [30], namely Chamfer Distance (CD) and F-score (F_A), with the latter aggregated as the area under the curve up to 1/20 of the datasets' maximum depth. All methods exploit GT intrinsics during evaluation. Nonetheless, we present results both with and without GT intrinsics for UniDepth.

Implementation Details. UniDepth is implemented in PyTorch [32] and CUDA [29]. For training, we use the AdamW [27] optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$) with an initial learning rate of 0.0001. The learning rate is divided by a factor of 10 for the backbone weights for every experiment and weight decay is set to 0.1. As the learning rate scheduler, we exploit Cosine Annealing to one-tenth starting from 30% of the training. We run 1M optimization iterations with a batch size of 128, each training dataset is uniformly represented in each batch. In particular, we sample 64 images and then we sample two different augmented views of the same image for consistency loss. The augmentations include both geometric and appearance (random brightness, gamma, saturation, hue shift, and grayscale) augmentations. ViT-L [11] backbone is initialized with weights from DINO-pre-

Table 1. **Comparison on zero-shot evaluation.** All methods are tested in a zero-shot setting on eight different datasets without overlap with any of the sets used for training. UniDepth-{C, V}: UniDepth-{ConvNext [26], ViT [11]}. (†): DDAD [19] in training set. (‡): predicted intrinsics are utilized for conditioning and backprojecting. Best viewed on a screen and zoomed in.

Method	NuScenes			DDAD			ETH3D			Diode (Indoor)			SUN-RGBD			VOID			IBims-1			HAMMER		
	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$
BTS [24]	33.7	68.0	37.5	43.0	40.8	40.5	26.8	29.9	27.4	19.2	22.8	31.6	76.1	14.6	64.8	47.4	25.8	64.5	53.1	17.5	57.2	3.89	20.9	22.8
AdaBins [2]	33.3	61.4	35.2	37.7	44.4	35.6	24.3	28.3	25.2	17.4	21.6	28.7	77.7	13.9	65.4	50.5	23.8	65.0	55.0	15.6	57.8	7.21	21.5	27.7
NeWCRF [50]	44.2	49.4	42.2	45.6	34.9	41.6	35.7	26.1	32.3	20.1	18.5	35.3	75.3	11.9	61.6	53.1	22.3	67.9	53.6	14.7	59.2	1.43	14.9	20.8
iDisc [34]	39.4	37.1	34.5	28.4	32.2	25.8	35.6	27.5	31.4	23.8	15.8	33.4	83.7	12.4	71.0	55.3	20.3	68.6	48.9	13.2	55.4	2.58	14.0	32.6
ZoeDepth [3]	28.3	31.5	26.0	27.2	31.7	21.1	35.0	17.6	26.4	36.9	12.8	40.5	86.7	9.58	75.6	63.4	15.9	72.4	58.0	10.9	59.6	0.72	9.78	21.0
Metric3D [†] [48]	72.3	29.0	53.9	—	—	—	45.6	18.9	35.9	39.2	11.1	42.1	15.4	13.4	14.4	65.9	16.2	70.4	79.7	10.1	68.5	3.40	12.1	29.0
UniDepth-C	83.3	22.9	62.3	83.2	21.4	59.3	49.8	13.2	33.7	60.2	9.03	50.0	94.8	8.10	81.4	86.6	12.8	85.1	79.7	8.92	66.7	20.2	8.78	57.1
UniDepth-V	86.2	21.7	64.2	86.4	20.3	61.8	32.6	11.6	24.3	77.1	6.38	59.4	96.6	7.05	81.9	89.4	10.9	85.7	23.9	7.22	37.1	13.3	7.41	55.9
UniDepth-C [‡]	83.3	22.9	60.9	83.1	21.4	57.3	22.9	13.1	25.4	60.4	9.01	49.9	92.3	8.27	75.2	86.5	12.8	85.0	79.4	8.88	64.2	12.7	9.30	54.8
UniDepth-V [‡]	86.2	21.7	63.0	86.4	20.3	60.4	17.6	11.4	21.4	77.4	6.36	58.6	94.8	7.17	75.9	90.2	10.9	86.2	17.5	7.20	36.5	2.56	8.35	53.8

Table 2. **Comparison on NYU test set.** The first five methods are trained on NYU and tested on it. The last four methods are tested in a zero-shot setting. UniDepth-{C, V}: UniDepth-{ConvNext [26], ViT [11]}. (†): MiDaS [35] pre-trained.

Method	$\delta_{0.5}$	δ_1	F_A	A.Rel	RMS	RMS _{log}	CD	SI_{log}
	Higher is better			Lower is better				
BTS [24]	66.1	88.5	74.0	10.9	0.391	0.141	0.160	11.5
AdaBins [2]	68.1	90.1	74.7	10.3	0.365	0.131	0.156	10.6
NeWCRF [50]	69.6	92.1	75.8	9.56	0.333	0.119	0.147	9.16
iDisc [34]	74.5	93.8	78.2	8.61	0.313	0.110	0.133	8.85
ZoeDepth [†] [3]	78.4	95.2	80.1	7.70	0.278	0.097	0.125	7.19
ZeroDepth [20]	—	90.1	—	10.0	0.380	—	—	—
Metric3D [48]	76.3	92.6	77.8	9.38	0.337	0.120	0.146	9.13
UniDepth-C	85.4	97.2	84.3	6.26	0.232	0.082	0.101	6.41
UniDepth-V	88.6	98.4	85.9	5.78	0.201	0.073	0.092	5.27

trained [5] models, and ConvNext-L [26] is ImageNet [8]-pre-trained. The required training time amounts to roughly 12 days on 8 NVIDIA A100. Ablations are conducted with three different seeds and for 100k training iterations, using a randomly sampled subset with a size equal to 20% of the original training set.

4.2. Comparison with the State of the Art

Our method consistently outperforms previous SotA methods as shown in Table 1. We particularly excel in the scale-invariant aspect, represented by SI_{log} , with an average 34.0% improvement, and an average 12.3% improvement for δ_1 and F_A . However, UniDepth could fail to capture the specific scene scales in certain cases, *e.g.* in ETH3D and IBims-1. This pitfall is demonstrated by the drop in scale-dependent metrics, *e.g.* F_A drop is 11.8% and 31.4%, respectively, although having a clear scale-invariant improvement of 36.9% and 28.5%. Therefore, we speculate that our method would still greatly benefit from domain-specific fine-tuning.

The last two rows in Table 1 present UniDepth in its whole design, namely functioning with solely the input image by self-prompting the predicted dense camera representation, as detailed in Eq. (3). Experiments show that not only is the performance preserved for most of the test sets, but UniDepth with the bootstrapped camera can also outperform models with GT camera, *e.g.* SI_{log} in ETH3D and IBims-1. On the other hand, in cases with particularly out-of-domain camera types, such as ETH3D or HAMMER, bootstrapping camera prediction results in additional noise for scaled depth prediction, thus worsening results for δ_1 .

Table 3. **Comparison on KITTI Eigen-split test set.** The first five methods are trained on KITTI and tested on it. The last four methods are tested in a zero-shot setting. UniDepth-{C, V}: UniDepth-{ConvNext [26], ViT [11]}. (†): MiDaS [35] pre-trained.

Method	$\delta_{0.5}$	δ_1	F_A	A.Rel	RMS	RMS _{log}	CD	SI_{log}
	Higher is better			Lower is better				
BTS [24]	86.9	96.2	82.0	5.63	2.43	0.089	0.42	8.18
AdaBins [2]	86.2	96.3	81.5	5.85	2.38	0.089	0.429	8.10
NeWCRF [50]	88.9	97.5	82.7	5.20	2.07	0.078	0.388	7.00
iDisc [34]	89.2	97.5	83.1	5.09	2.07	0.077	0.380	7.11
ZoeDepth [†] [3]	87.4	96.5	82.1	5.76	2.39	0.089	0.431	7.47
ZeroDepth [20]	—	89.2	—	10.2	4.38	0.196	—	—
Metric3D [48]	88.9	97.5	82.9	5.33	2.26	0.081	0.392	7.28
UniDepth-C	91.1	97.9	83.9	4.69	2.00	0.072	0.371	6.71
UniDepth-V	93.4	98.6	85.0	4.21	1.75	0.064	0.338	5.84

Table 2 and Table 3 display results on the two popular benchmark NYU [28] and KITTI [17] Eigen-split. UniDepth sets the state of the art in these two benchmarks despite being compared with models trained on the same domain. Importantly, the KITTI Depth Prediction Benchmark, which provides a perfectly fair evaluation, underscores the excellent zero-shot performance of our method and its robustness compared to the current MMDE SotA methods, as UniDepth ranks first on this benchmark at the time of submission, with a 15.5% improvement in SI_{log} over the second-best method. Performance disparities are not solely attributed to dataset characteristics, as observed in the comparison with Metric3D and ZeroDepth. Despite being trained on a smaller dataset, UniDepth outperforms both of these methods. In particular, UniDepth improves in δ_1 over Metric3D and ZeroDepth by 5.8% and 7.3%, respectively, on NYU (Table 2) and by 1.1% and 9.4%, respectively, on KITTI (Table 3). Moreover, ZoeDepth, which has a capacity similar to our ViT-based approach and is pre-trained on the diverse MiDaS dataset [35], shows limitations in general zero-shot scenarios in Table 1, exhibiting performance comparable to traditional MMDE methods especially on scale-invariant metrics.

For the sake of fair comparison, we provide in Table 4 a comparison between Metric3D, iDisc, and UniDepth where the latter two are retrained on a *strict* subset of Metric3D’s data, namely accounting for one-quarter of the original Metric3D dataset, with same framework detailed in Sec. 4.1. The results are two-fold: they demonstrate how UniDepth still surpasses Metric3D with a subsplit of the training set, and how MMDE SotA methods designed for single-domain can

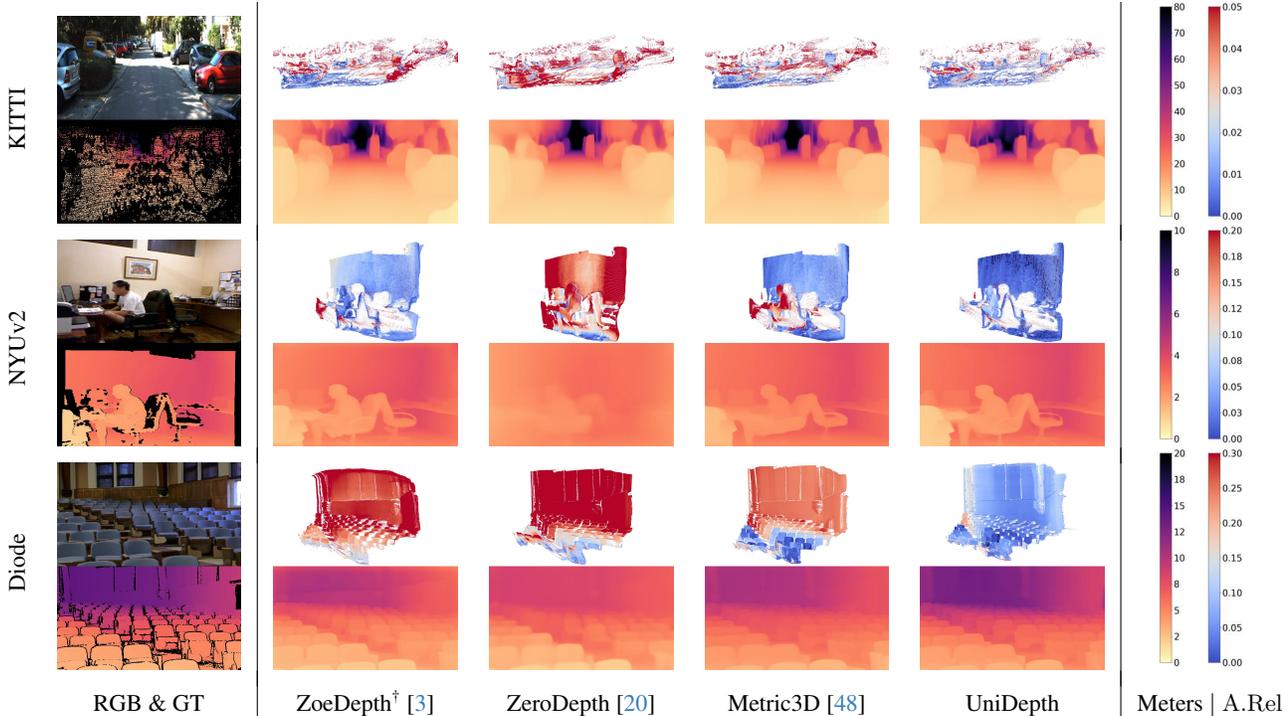


Figure 4. **Zero-shot qualitative results.** Each pair of consecutive rows corresponds to one test sample. Each odd row shows the input RGB image and the predicted pointcloud color-coded with *coolwarm* based on the absolute relative error. Each even row shows GT depth and the predicted depth. The last column represents the specific colormap ranges for depth and error. (†): KITTI and NYU in the training set.

Table 4. **Comparison with equivalent training setup.** All methods have the same backbone, ConvNext-L [26] and are tested in a zero-shot regime on KITTI Eigen-split and NYU. iDisc and UniDepth are retrained on a strict subset of Metric3D for 500k iterations as in [48].

Method	KITTI			NYU		
	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$	$\delta_1 \uparrow$	$SI_{log} \downarrow$	$F_A \uparrow$
iDisc [34]	93.4	8.36	78.0	92.1	<u>8.82</u>	75.0
Metric3D [48]	<u>97.5</u>	<u>7.28</u>	<u>82.9</u>	<u>92.6</u>	9.13	<u>77.8</u>
UniDepth	97.9	6.66	83.8	97.1	6.69	84.3

not fully exploit the training diversity. Qualitative results in Fig. 4 emphasize how the method excels in capturing the overall scale and scene complexity in a zero-shot setup.

4.3. Ablation Study

The importance of each component introduced in UniDepth in Sec. 3 is evaluated by ablating the method in Table 5. All ablations exploit the predicted camera representation, if not stated otherwise. The first distinction involves the *Oracle* model, which operates under ideal conditions with known camera information during training and testing, addressing a task similar to [20, 48]. On the other hand, *Baseline* is a straightforward encoder-decoder implementation with a (x,y,z) output, as outlined at the beginning of Sec. 3, while *Baseline++* exploits the proposed pseudo-spherical representation. Modules’ architectures are consistent across experiments. The *In-Domain* column reflects testing on valida-

tion splits of training domains, while *Out-of-Domain* corresponds to zero-shot testing, as detailed in Sec. 4.1. Notably, *In-Domain* results exhibit a higher degree of homogeneity compared to *Out-of-Domain*, which is noisier yet more informative for gauging expected performances in downstream applications and in-the-wild deployment.

Architecture. The *Oracle* model demonstrates more robust scale-dependent performance during zero-shot testing compared to the *Full* model, highlighting how the proposed task is inherently more demanding. The *Baseline* model illustrates an approach to the problem without utilizing external information and lacking a proper design for both internal and output space. This approach yields markedly inferior results for both *In-Domain* and *Out-of-Domain* scenarios in terms of depth and 3D reconstruction metrics.

Camera Module. In Table 5, row 3, the benefit of the Camera Module becomes apparent, revealing a substantial disparity in the effect of this module on scale-invariant and scale-dependent metrics for in- and out-of-domain testing. This disparity stems from the absence of prior knowledge of the model regarding scale, impeding its optimal utilization of the diverse training set. Concentrating solely on predicting depth, rather than a complete 3D output, proves advantageous in averting convergence issues during training. This is evident in comparison with methods predicting 3D, either without reliance on camera information (rows 8 and 9) or influenced by intertwined optimization (row 4), as elu-

Table 5. **Ablations of UniDepth.** *In-Domain* corresponds to the union of the training domain’s validation sets, while *Out-of-Domain* involves the union of zero-shot testing sets. *Oracle* is the model with provided GT cameras at training and test time. *Baseline* directly predicts 3D points in Cartesian space, *Baseline++* in pseudo-spherical. *Full* represents the final UniDepth. All models have the same depth and camera module architecture, if any. ARel_C is the mean of elementwise absolute relative error for camera intrinsics. (\dagger): GT camera intrinsics utilized for backprojection. The backbone used is ConvNext-L [26]. Medians and median average deviations over three runs are reported.

Ablation	In-Domain				Out-of-Domain			
	$\delta_1 \uparrow$	$\text{SI}_{\log} \downarrow$	$F_A \uparrow$	$\text{ARel}_C \downarrow$	$\delta_1 \uparrow$	$\text{SI}_{\log} \downarrow$	$F_A \uparrow$	$\text{ARel}_C \downarrow$
1 Oracle \dagger	89.06 \pm 0.03	13.15 \pm 0.02	65.45 \pm 0.13	n/a	68.11 \pm 0.17	14.78 \pm 0.01	57.17 \pm 0.09	n/a
2 Full	88.89 \pm 0.10	13.13 \pm 0.01	63.52 \pm 0.08	2.05 \pm 0.01	57.06 \pm 1.48	14.83 \pm 0.04	49.71 \pm 0.55	13.54 \pm 0.85
3 – Camera \dagger	87.42 \pm 0.04	13.49 \pm 0.08	63.78 \pm 0.02	n/a	48.38 \pm 0.97	15.55 \pm 0.15	45.21 \pm 0.86	n/a
4 – Spherical	61.30 \pm 1.00	19.36 \pm 0.09	17.89 \pm 0.11	48.29 \pm 4.03	37.09 \pm 1.37	22.49 \pm 0.16	21.78 \pm 0.14	87.51 \pm 11.1
5 – \mathcal{L}_{con}	88.53 \pm 0.07	13.24 \pm 0.01	60.89 \pm 0.15	2.65 \pm 0.06	52.89 \pm 0.21	14.85 \pm 0.01	45.17 \pm 0.32	14.27 \pm 0.41
6 – Dense	87.62 \pm 0.11	13.41 \pm 0.05	61.33 \pm 0.54	1.91 \pm 0.04	55.65 \pm 0.18	15.04 \pm 0.04	43.19 \pm 0.24	16.61 \pm 0.41
7 – Detach	88.16 \pm 0.12	13.48 \pm 0.06	64.19 \pm 0.17	0.93 \pm 0.02	46.60 \pm 0.25	15.26 \pm 0.10	43.85 \pm 2.01	18.99 \pm 1.00
8 Baseline	77.36 \pm 0.22	21.17 \pm 0.28	16.29 \pm 0.26	n/a	48.19 \pm 1.02	23.05 \pm 0.45	14.29 \pm 0.36	n/a
9 Baseline++	82.41 \pm 0.13	16.31 \pm 0.05	41.98 \pm 0.12	n/a	51.22 \pm 0.35	18.14 \pm 0.05	38.27 \pm 0.02	n/a

culated in Sec. 3. Refraining from relying on the camera also constrains the model’s capacity to recover a multimodal distribution for out-of-domain samples. The lack of a (bootstrapped) prior prevents the depth module from serving as a corrective mechanism based on an initial scale estimation and imposes an unnecessary computational burden, *i.e.* recovering the depth values from scratch. This limitation is underscored by the marked variability observed for test sets strongly out-of-distribution, such as KITTI, when comparing the utilization or absence of camera information (rows 2 and 3, respectively). In particular, *Full* achieves 95.2% in δ_1 in KITTI, while “– Camera” obtains 58.9% for the same test set, despite a mere 2% difference between the two versions on nuScenes and DDAD.

Optimization and Output Representation. All ablations employ the same loss $\mathcal{L}_{\lambda \text{MSE}}$, but across different output spaces. In row 4, a Cartesian output space is used instead of a pseudo-spherical from Sec. 3.1, which results in substantially inferior performance due to the respective intertwined formulation of camera and depth output spaces. The *Baseline* (row 8) also employs a Cartesian representation, but the negative impact of this choice is less pronounced in this model because of the absence of a camera module. More specifically, the decoder of *Baseline* is not conditioned on inaccurate prior camera and scale information as in row 4. Moreover, row 9 corresponds to *Baseline* with pseudo-spherical representation. Comparison between row 8 and row 9 shows that when predicting directly the 3D outputs, the choice of the output representation is still relevant in defining a better internal representation and optimization. Row 5 demonstrates the positive impact of the geometric invariance loss. This loss contributes to enhanced in-domain and out-of-domain performance by promoting the invariance of depth features to appearance variations owing to different camera intrinsics. Furthermore, stopping the gradient from propagating from the Camera Module to the Encoder (row 7), as described in Sec. 3.4, proves particularly beneficial in avoiding scale

and camera overfitting in zero-shot testing, and stabilizes the training. The more stable training is obtained by limiting the dominant effect that camera supervision has on the gradient of the Encoder weights compared to depth supervision.

Camera Representation. In row 6, the model incorporates a sparse camera representation, specifically the pinhole camera model with (f_x, f_y, c_x, c_y) , leading to sparse camera prompting and scalar supervision; the camera module still predicts the residual components as outlined in Sec. 3.2. This approach hurts generalization, as evidenced by ARel_C in the out-of-domain evaluation, despite the slight improvement in in-domain ARel_C . We speculate that the four prompts convey less robust information to the depth module than their dense counterpart, resulting in inferior performance for depth metrics compared to *Full* for both in- and out-of-domain.

5. Conclusion

In this work, we propose UniDepth to predict metric 3D points in diverse scenes relying solely on a single input image. Through meticulous ablation studies, we systematically address the challenges inherent in universal MMDE tasks, underscoring the pivotal contributions of our work. The designed self-prompting camera allows camera-free test time application and renders the model more robust against camera noise. The introduced pseudo-spherical output space representation adequately disentangles the camera and depth of the optimization process. Furthermore, the proposed geometric invariance loss effectively ensures camera-aware depth consistency. Extensive validations unequivocally exhibit how UniDepth sets the new state of the art across multiple benchmarks in a zero-shot regime, even surpassing in-domain trained methods. This attests to the robustness and efficacy of our model and, most importantly, outlines its potential to propel the field of MMDE to new frontiers.

Acknowledgment. This work is funded by Toyota Motor Europe via the research project TRACE-Zürich.

References

- [1] Manuel Lopez Antequera, Pau Gargallo, Markus Hofinger, Samuel Rota Bulò, Yubin Kuang, and Peter Kotschieder. Mapillary planet-scale depth dataset. In *The European Conference Computer Vision (ECCV)*, pages 589–604. Springer International Publishing, 2020. 2, 5
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4008–4017, 2020. 1, 2, 6
- [3] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 2, 6, 7
- [4] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9650–9660, 2021. 6
- [6] 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 (CVPR)*, 2016. 5
- [7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. 6
- [9] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. 1
- [10] Xingshuai Dong, Matthew A Garratt, Sreenatha G Anavatti, and Hussein A Abbas. Towards real-time monocular depth estimation for robotics: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):16940–16961, 2022. 1
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*. OpenReview.net, 2021. 4, 5, 6
- [12] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10786–10796, 2021. 2
- [13] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. pages 2366–2374. Neural information processing systems foundation, 2014. 1, 2, 5
- [14] Jose M Facil, Benjamin Ummenhofer, Huizhong Zhou, Luis Montesano, Thomas Brox, and Javier Civera. Cam-convs: Camera-aware multi-scale convolutions for single-view depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11826–11835, 2019. 2
- [15] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2002–2011, 2018. 1, 2
- [16] Ravi Garg, B. G. Vijay Kumar, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. *Lecture Notes in Computer Science*, 9912 LNCS:740–756, 2016. 5
- [17] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 2, 5, 6
- [18] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S. Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, Tiffany Fernandez, Martin Jänicke, Sudesh Mirashi, Chiragkumar Savani, Martin Sturm, Oleksandr Vorobiov, Martin Oelker, Sebastian Garreis, and Peter Schuberth. A2D2: Audi Autonomous Driving Dataset. *arXiv preprint arXiv:2004.06320*, 2020. 5
- [19] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5, 6
- [20] Vitor Guizilini, Igor Vasiljevic, Dian Chen, Rares Ambrus, and Adrien Gaidon. Towards zero-shot scale-aware monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9233–9243, 2023. 1, 2, 5, 6, 7
- [21] HyunJun Jung, Patrick Ruhkamp, Guangyao Zhai, Nikolas Brasch, Yitong Li, Yannick Verdie, Jifei Song, Yiren Zhou, Anil Armagan, Slobodan Ilic, et al. Is my depth ground-truth good enough? HAMMER – Highly Accurate Multi-Modal dataset for dEnse 3D scene Regression. *arXiv preprint arXiv:2205.04565*, 2022. 5
- [22] Tobias Koch, Lukas Liebel, Marco Körner, and Friedrich Fraundorfer. Comparison of monocular depth estimation methods using geometrically relevant metrics on the IBims-1 dataset. *Computer Vision and Image Understanding (CVIU)*, 191:102877, 2020. 5

- [23] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. *Proceedings of the International Conference on 3D Vision (3DV)*, pages 239–248, 2016. [2](#)
- [24] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *CoRR*, abs/1907.10326, 2019. [2](#), [5](#), [6](#)
- [25] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 38:2024–2039, 2015. [2](#)
- [26] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11976–11986, 2022. [6](#), [7](#), [8](#)
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *7th International Conference on Learning Representations, ICLR 2019*, 2017. [5](#)
- [28] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *The European Conference Computer Vision (ECCV)*, 2012. [2](#), [5](#), [6](#)
- [29] John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. Scalable parallel programming with cuda: Is cuda the parallel programming model that application developers have been waiting for? *Queue*, 6(2):40–53, 2008. [5](#)
- [30] Evin Pinar Örnek, Shristi Mudgal, Johanna Wald, Yida Wang, Nassir Navab, and Federico Tombari. From 2d to 3d: Rethinking benchmarking of monocular depth prediction. *arXiv preprint arXiv:2203.08122*, 2022. [5](#)
- [31] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [1](#)
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8024–8035. Curran Associates, Inc., 2019. [5](#)
- [33] Vaishakh Patil, Christos Sakaridis, Alexander Liniger, and Luc Van Gool. P3Depth: Monocular depth estimation with a piecewise planarity prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1600–1611. IEEE, 2022. [1](#), [2](#)
- [34] Luigi Piccinelli, Christos Sakaridis, and Fisher Yu. iDisc: Internal discretization for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#), [2](#), [5](#), [6](#), [7](#)
- [35] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 44(3):1623–1637, 2020. [2](#), [6](#)
- [36] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12159–12168, 2021. [1](#)
- [37] Thomas Schöps, Johannes L. Schönberger, 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [5](#)
- [38] Shuran Song, Samuel P. Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 07-12-June-2015:567–576, 2015. [5](#)
- [39] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2446–2454, 2020. [5](#)
- [40] Igor Vasiljevic, Nicholas I. Kolkin, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z. Dai, Andrea F. Daniele, Mohammadreza Mostajabi, Steven Basart, Matthew R. Walter, and Gregory Shakhnarovich. DIODE: A dense indoor and outdoor depth dataset. *CoRR*, abs/1908.00463, 2019. [5](#)
- [41] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudolidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8445–8453, 2019. [1](#)
- [42] Yan Wang, Xiangyu Chen, Yurong You, Li Erran Li, Bharath Hariharan, Mark Campbell, Kilian Q. Weinberger, and Wei Lun Chao. Train in germany, test in the usa: Making 3d object detectors generalize. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11710–11720, 2020. [1](#)
- [43] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Advances in Neural Information Processing Systems*, 2021. [5](#)
- [44] Alex Wong, Xiaohan Fei, Stephanie Tsuei, and Stefano Soatto. Unsupervised depth completion from visual inertial odometry. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):1899–1906, 2020. [5](#)
- [45] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [5](#)

- [46] Guanglei Yang, Hao Tang, Mingli Ding, Nicu Sebe, and Elisa Ricci. Transformer-based attention networks for continuous pixel-wise prediction. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16249–16259, 2021. [2](#)
- [47] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 204–213, 2021. [2](#)
- [48] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9043–9053, 2023. [1](#), [2](#), [5](#), [6](#), [7](#)
- [49] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multi-task learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2636–2645, 2020. [5](#)
- [50] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. Neural window fully-connected crfs for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3906–3915. IEEE, 2022. [1](#), [2](#), [6](#)
- [51] Amir R Zamir, Alexander Sax, William B Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2018. [5](#)
- [52] Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun. Does computer vision matter for action? *Science Robotics*, 4, 2019. [1](#)