

Extending Foundational Monocular Depth Estimators to Fisheye Cameras with Calibration Tokens

Suchisrit Gangopadhyay^{1*}, Jung-Hee Kim^{2*}, Xien Chen^{1*}, Patrick Rim¹, Hyoungseob Park¹, Alex Wong¹
Yale University¹, Michigan State University²

{rit.gangopadhyay,xien.chen,patrick.rim,hyoungseob.park,alex.wong}@yale.edu, kimjun84@msu.edu

Abstract

We propose a method to extend foundational monocular depth estimators (FMDEs), trained on perspective images, to fisheye images. Despite being trained on tens of millions of images, FMDEs are susceptible to the covariate shift introduced by changes in camera calibration (intrinsic, distortion) parameters, leading to erroneous depth estimates. Our method aligns the distribution of latent embeddings encoding fisheye images to those of perspective images, enabling the reuse of FMDEs for fisheye cameras without retraining or finetuning. To this end, we introduce a set of Calibration Tokens as a light-weight adaptation mechanism that modulates the latent embeddings for alignment. By exploiting the already expressive latent space of FMDEs, we posit that modulating their embeddings avoids the negative impact of artifacts and loss introduced in conventional recalibration or map projection to a canonical reference frame in the image space. Our method is self-supervised and does not require fisheye images but leverages publicly available large-scale perspective image datasets. This is done by recalibrating perspective images to fisheye images, and enforcing consistency between their estimates during training. We evaluate our approach with several FMDEs, on both indoors and outdoors, where we consistently improve over state-of-the-art methods using a single set of tokens for both. Code available at: github.com/JungHeeKim29/calibration-token.

1. Introduction

Three-dimensional (3D) reconstruction is a fundamental component in many spatial applications, including autonomous vehicles, extended reality (XR), robotic manipulation. Each of these applications has unique demands for the field of view (FOV), often wider than the standard (perspective) camera. To meet this need, these applications tend to be deployed on systems equipped with fish-




Figure 1. **Foundational monocular depth estimators fail on fisheye images.** Despite being trained on large-scale datasets, foundational monocular depth estimators (FMDEs) models produces erroneous outputs. The inaccurate, blurry estimates are caused by a covariate shift that stem from fisheye distortion.

eye or other wide-angle cameras, which allows for wider coverage of the 3D environment. However, images captured by these cameras also come with substantial distortion, which arise from differences in projective geometry, where straight lines within the 3D environment or the 3D scene are preserved in perspective images but may appear curved in fisheye images.

Foundational monocular depth estimators (FMDEs) [44, 48, 73] are trained on orders of tens of millions of images, enabling them to generalize across a wide range 3D scenes. However, their training data is comprised of internet images, which are predominantly captured using perspective cameras. Hence, despite being trained on large-scale datasets, FMDEs produce erroneous estimates when transferred to fisheye images (see Fig. 1). These errors stem from a covariate shift, which can be characterized by changes in camera calibration (intrinsic, distortion) parameters – leading to differences in object appearance and their perceived depth or distance from the camera.

To address fisheye distortion, one solution is to recalibrate and undistort images or perform a map projection to some canonical reference frame. In principle, if one has the correct calibration, it is possible to re-project a fisheye image into a perspective-like view (or vice versa). In practice, however, there are several problems: (1) The calibration process itself can be error-prone and sensitive to physical perturbations in the camera system. Minor bumps, focus changes, or lens replacements can degrade or inval-

*Equal contribution

update previously computed intrinsic parameters. (2) Even when re-projection is performed accurately, the transformation introduces latency and spatial artifacts (e.g., stretching, cropping, aliasing, loss). When used as a preprocessing step for existing pretrained depth estimators, these artifacts still present a covariate shift and can degrade performance.

Another solution is to train a separate monocular depth estimator specifically for fisheye images. However, publicly available image datasets for fisheye cameras are orders of tens to hundreds of times smaller than those for perspective cameras. Hence, it is difficult to assemble sufficient data to reach the large-scale training requirement of an FMDE. Nonetheless, one can adapt or finetune existing FMDEs for fisheye imagery. While this can improve performance on fisheye images, it introduces the risk of parameter drift, where the resulting FMDEs may lose their generalizability across 3D scenes. Moreover, the resulting finetuned model becomes specialized to fisheye cameras, limiting its applicability to other camera types, which adds operational overhead in applications involving mixed camera systems, e.g., autonomous vehicles or robotics.

To address these challenges, we propose a novel approach termed *Calibration Tokens*. Our key insight is that existing FMDEs are already capable of estimating depth for perspective images, and that errors on fisheye images are caused by a covariate shift due to differences in camera calibration and distortion. Hence, rather than retraining or finetuning the entire model, we aim to “recalibrate” the fisheye latent embeddings such that they become more conducive to an FMDE originally trained on perspective images. Leveraging the fact that many FMDEs [44, 48, 49, 73] follow a Transformer-based architecture [11], we will exploit the (self- and cross-) attention mechanism to modulate the latent (token) embeddings by inserting Calibration Tokens as part of the input. Therefore, the existing FMDE will remain effectively unchanged, while Calibration Tokens serve to adapt their internal representations to mitigate the covariate shift by aligning the latent embeddings of fisheye images to the distribution of latent embeddings of perspective images. This design also allows us to preserve the original image content without performing any spatial re-projection, ensuring the process is lossless in terms of the raw pixels. Our hypothesis is that by adding a small set of trainable tokens to encode the fisheye camera calibration information and utilizing them to recalibrate the latent embeddings, we will be able to reuse existing FMDEs trained on perspective images and adapt them to fisheye images without sacrificing their generalizability across diverse 3D scenes.

To train these Calibration Tokens, we propose a self-supervised objective that leverages inverse warping in the input and output spaces. FMDEs can infer high-fidelity depth maps for perspective images, so we use the perspective image depth estimates as our training target. We then

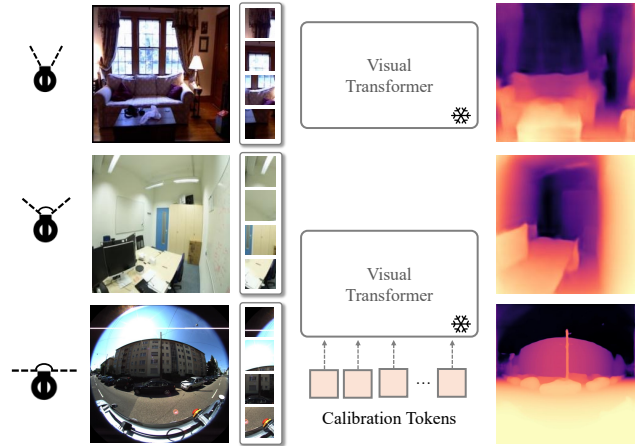


Figure 2. **Inference on different cameras.** Calibration Tokens enable foundational monocular depth estimators to adapt to fisheye images while maintaining performance on perspective images.

induce artificial distortion on perspective images to create pairs of perspective and synthetic fisheye images with diverse fisheye distortions. However, rather than doing the same in the output space, we undistort the fisheye depth maps to the original perspective reference frame to compute a self-supervised loss between the undistorted fisheye and perspective depth maps. By minimizing this self-supervised loss, the Calibration Tokens learn to align fisheye image embeddings to those of perspective images in the latent space, without any labels. Additionally, computing the loss in the original perspective frame allows our method to preserve the supervision signal instead of introducing artifacts.

Our approach allows us to bypass the need to compile large-scale fisheye datasets by exploiting the abundance of perspective image datasets. As our method operates in the reference frame of the input, we also avoid transformation artifacts at inference time, whether in the input or output space. Furthermore, our method preserves compatibility with perspective images: One simply needs to append or remove Calibration Tokens for FMDEs to be applied to fisheye or perspective images. We demonstrate our method on indoors and outdoors across several recent FMDEs and consistently improve over baselines.

Our contributions: (1) We propose a novel approach to extend foundational monocular depth estimators (FMDEs) trained on perspective images to fisheye images. (2) We introduce Calibration Tokens that modulate the latent embeddings of fisheye images towards the distribution of perspective image embeddings. (3) We introduce a self-supervised training objective that recalibrates input perspective image to fisheye images, but “undos” the transformation in the output to enable loss computation on high-fidelity (perspective) depth maps inferred by FMDEs. (4) Our approach only requires training one set of tokens to achieve state-of-the-art performance for both indoors and outdoors.

2. Related Works

Monocular Depth Estimation can be trained in a supervised or unsupervised manner. *Supervised methods* [13, 14, 29–32, 37, 71, 77] minimize the difference between depth estimates and ground-truth depth maps. [18] re-formulated the problem as ordinal regression while other methods proposed architectures innovations. [2] partitions depth ranges into adaptive bins. [5] incorporates an attention-based up-sample block. [33] employs hierarchical aggregation and heterogeneous interaction modules. [80] uses neural window fully-connected CRFs to compute energy. [59] synthesizes perspectively accurate images to enrich training data. Additional inputs e.g., language [81–83], lidar [6, 15, 51, 74], radar [50, 54], are used to enable metric-scale depth estimates. *Unsupervised methods* [7, 8, 23, 38, 42, 47, 57, 62–67, 84] minimize photometric reconstruction error. [19] frames depth estimation as a novel view synthesis problem. [20] introduces a left-right consistency loss. [88] uses a pose network to enable unsupervised training on video sequences. [21] introduced auto-masking and min-reprojection loss. Additional loss terms based on visual odometry [16, 60], iterative closest point [40], surface normals [75], trinocular assumption [46], and semantic segmentation [22, 28] were also introduced. [39] redesigned the skip connection and decoders to extract high-resolution features, [86] combined global and local representations and [85] introduced a lightweight architecture with dilated convolution and attention. AugUndo [68] leveraged invertibility of transformation groups for data augmentation.

Foundational Monocular Depth Estimators are trained with supervised or semi-supervised learning on large-scale datasets. MiDaS [49] is the first to demonstrate generalizable monocular depth estimation by compiling datasets for large-scale training. DPT [48] extended the approach and introduced transformers for dense predictions. Marigold [26] repurposes diffusion models for monocular depth estimation. DepthAnything [73] proposes a pseudo-labeling method to curate a large-scale dataset. Additionally, UniDepth [44] employs a camera self-prompting module and a pseudo-spherical output space, enabling metric-scale depth prediction across diverse 3D scenes without relying on external camera parameters. DepthPro [3] proposes a multi-scale vision transformer for metric-scale depth estimation. As all of these FMDEs are trained on perspective images, they fail to generalize to fisheye cameras.

Fisheye Images. Images taken by a fisheye camera are distorted and unsuitable for use in a perspective image encoder. Existing distortion correction algorithms [12, 25] rely on different camera projection models [27, 41, 55] to undistort images into a perspective view. However, these methods depend on camera calibration parameters, which can introduce artifacts due to calibration inaccuracies. Recent approaches

[24, 36] demonstrate training a separate model to perform depth estimation with different camera types. They utilize an equirectangular projection to project points from different reference frames to a canonical equirectangular frame, but this can incur transformation artifacts and distortions. Additionally, deep-learning-based methods [17, 34] that aim to rectify distortion have been introduced. However, these methods require a large number of parameters with limited accuracy and field of view. Consequently, many recent works targeted for fisheye images involve training an entire network [1, 79, 87] exclusively on fisheye images. Our method extends foundational monocular depth estimators to fisheye images instead.

Token-Based Methods. Recent transformer-based architectures represent input images [11] (or other modalities [69, 70]) as sequences of tokens. In many cases, an additional token (e.g. [CLS] token in BERT [10] or the distillation token in DeiT [58]) is employed to aggregate information across all tokens. Such tokens can be adapted to various purposes, acting as a compact representation that “binds” or fuses information, e.g., [72] uses tokens to learn synthetic and real tactile response maps. Inspired by these advances, we introduce minimal trainable tokens appended to the fisheye embedding, enabling the model to “bind” or reconcile image distortions within a frozen backbone. Our approach is lightweight and requires no major architectural modifications, but extends foundational monocular depth estimators trained on perspective images to fisheye images.

3. Method

Let $I : \Omega \mapsto \mathbb{R}^3$ denote an RGB image obtained from a calibrated camera and $\Omega \subset \mathbb{R}^2$ the image space. Monocular depth estimation aims to learn a parameterized function $h_{\omega, \psi} : \mathbb{R}^3 \rightarrow \mathbb{R}_+$ that maps an image to a depth map $d : \Omega \mapsto \mathbb{R}_+$. We assume access to a foundational monocular depth estimator (FMDE) pretrained on some large-scale dataset of perspective images. We will pair each image with (pseudo)ground truth $\tilde{d} = h_{\omega, \psi}(I)$ inferred by the FMDE to obtain a training dataset $\mathcal{D} = \{(I^{(n)}, \tilde{d}^{(n)})\}_{n=1}^N$.

To extend FMDEs, trained on perspective images, to fisheye images, we introduce Calibration Tokens as an adaptation mechanism. Due to the prevalence of Transformer architectures in many FMDEs, we train a set of lightweight tokens to model the change in calibration between a perspective camera and different fisheye cameras. The goal of our Calibration Token is to recalibrate or translate latent embeddings of fisheye images back to those of perspective images. Our method takes advantage of the attention mechanism inherent in FMDEs and enable Calibration Tokens to modulate the latent embeddings, thus facilitating latent alignment. The outcome is an FMDE that is capable of inferring depth for fisheye images with Calibration Tokens, and perspective images without.

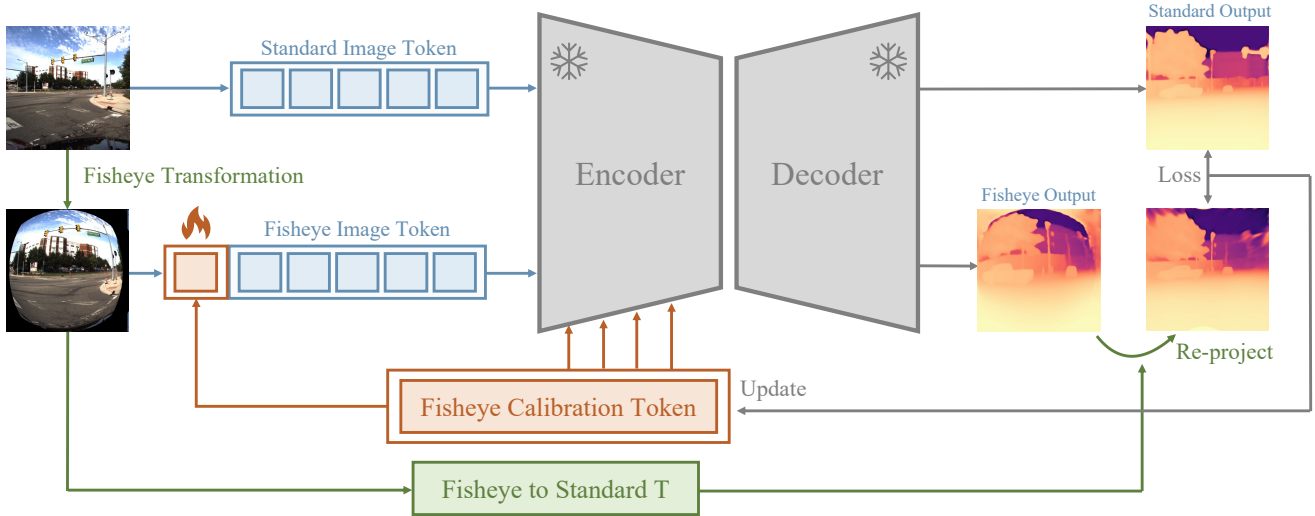


Figure 3. **Overview of our method.** We introduce a set of trainable *Calibration Tokens*, which is appended to the input sequence of the fisheye image tokens. The Calibration Tokens are trained to adapt the model to produce accurate depth maps for images with various fisheye distortions. A unique fisheye calibration token is appended to the input of each new layer of the encoder.

3.1. Extending FMDEs with Calibration Tokens

Specifically, let $\phi \in \mathbb{R}^{M \times F}$ represent a set of Calibration Tokens, where M denotes the number of tokens and F their number of dimensions. For a given layer within the encoder f_ω of an FMDE, we will concatenate Calibration Tokens to the input sequence of patches or (embeddings) of the vision transformer: $f_\omega([I; \phi]) = [z^{(L)}; f_\omega(\phi)]$, where z denotes the latent embeddings recalibrated by ϕ , L the last layer, and $[\cdot; \cdot]$ the concatenation operation.

As each layer denotes a separate latent space, we extend our approach to a multi-layer modulation scheme. Let $\Phi \in \mathbb{R}^{L \times M \times F}$ be the set of Calibration Tokens for each layer $l \in \{1, \dots, L\}$. A unique set of Calibration Tokens $\phi^{(l)} \in \mathbb{R}^{M \times F}$ is appended at each encoder layer: $f_\omega^{(l)}([z^{(l-1)}; \phi^{(l)}]) = [z^{(l)}; f_\omega^{(l)}(\phi^{(l)})]$ for a layer l . Each set of Calibration Tokens will modulate the patch embeddings for a specific layer through the attention mechanism; hence, following the convention in existing works [4, 9], we discard Calibration Tokens from the encoder output. A key insight is that the FMDE is already able to estimate high-fidelity depth maps for perspective images. We posit that the covariate shift exist in the encodings of fisheye image. Hence, once the latent embeddings of fisheye images have been recalibrated to those of perspective images, the decoder will be able to estimate depth to similar fidelity as perspective images. Therefore, we do not utilize Calibration Tokens to modulate the decoder layers. The final estimate is obtained by $\hat{d} = g_\psi(z^{(L)})$, where g_ψ denotes the decoder.

Since our method does not apply spatial transformations during inference, it remains entirely *lossless* for input images. Additionally, it offers several efficiencies: (1) The

only trainable parameters in our method are the light-weight Calibration Tokens, which consist of significantly fewer parameters than vision transformer models. Our method introduces minimal computational overhead and results in lower time and space complexity than training or finetuning a full model. (2) Our approach is backward-compatible with perspective images. By omitting our Calibration Tokens, an FMDE maintains its original depth estimation performance on perspective images. (3) At inference, camera intrinsics are not required, as the training process allows generalization across various fisheye camera intrinsics. As a result, our method eliminates the need for the arduous and error-prone calibration process after training.

3.2. Learning Calibration Tokens

To train Calibration Tokens, we will leverage the abundance of publicly available perspective image datasets. During our training, we synthesize fisheye images from perspective images by recalibrating them using artificial fisheye intrinsic and distortion parameters. This will produce pairs of perspective and fisheye images from which we can leverage self-supervision, and allows us to use a much larger training dataset than exclusively training with real fisheye images. We follow previous approaches [17, 78] to obtain synthetic fisheye images from the calibrated perspective images. Our synthesized fisheye images follow the distortion model introduced by Kannala & Brandt [25]:

$$r(\theta) = k_1\theta + k_2\theta^3 + k_3\theta^5 + k_4\theta^7, \quad (1)$$

where θ denotes the angle between the ray and the optical axis, and $\{k_i\}_{i=1}^4$ are distortion coefficients that can represent most of the real world fisheye distortion models. The

change in coordinate between (x, y) in the perspective image and (x', y') in the fisheye image can be formulated as

$$\begin{aligned} x' &= r(\theta) \cos(\varphi), \quad y' = r(\theta) \sin(\varphi), \\ \varphi &= \arctan((y - c_x)/(x - c_y)), \end{aligned} \quad (2)$$

where (c_x, c_y) is the principal point in the perspective image. We define the transformation from the perspective to the fisheye reference frame as T and its inverse transformation as T^{-1} . Our training dataset is composed of perspective images and synthesized distorted image pairs with corresponding forward and inverse transformations.

Loss Function. Inspired by AugUndo [68] and their use of invertible transformations to preserve the supervision signal by undoing data augmentation, we propose to synthesize fisheye images from the abundance of perspective images as inputs, but undistort the output to facilitate loss computation. By remapping depth estimates of synthetic fisheye images to the perspective frame, we enable the use of high-fidelity estimates inferred by FMDEs on perspective images as supervision. Hence, we can optimize Calibration Tokens with the following self-supervised loss:

$$\arg \min_{\Psi} \frac{1}{N} \sum_{n=1}^N \sum_{x \in \Omega} \log(|\tilde{d}^{(n)}(x) - T^{-1} \circ \hat{d}^{(n)}(x)| + 1), \quad (3)$$

where $\tilde{d} = h_{\omega, \psi}(I)$ and $\hat{d} = h_{\omega, \psi}(T \circ I; \Phi)$ denotes the predicted depth map from the given perspective image and synthesized fisheye image, respectively. Φ denotes the proposed trainable Calibration Token appended to the patch embeddings. Calibration Tokens are trained to minimize the difference between the perspective output and the fisheye output re-projected into the perspective reference frame. Eq. (3) follows the log of absolute differences (logL1) proposed in [43], which enhances training stability and empirically outperforms L1 loss, especially in border regions where discrepancies between perspective and fisheye images are most significant (see Sec. 4.3 for details).

It is important to note that attempting to instead transform the perspective depth maps outputted by FMDEs to the fisheye reference frame for the loss computation introduces information loss in the training objective. This will lead to re-projection artifacts in the supervision and result in learning inaccuracies during training. In Section A of the Supp. Mat., we further demonstrate comparison results between training in fisheye image space and perspective image space. Our training scheme is self-supervised and requires only calibrated perspective images, which can be easily obtained, making our approach both scalable and practical.

4. Experiments

Datasets. Training our Calibration Tokens requires only calibrated perspective images, enabling us to leverage significantly more data compared to training solely on fisheye

images. Moreover, since our loss is computed based on comparisons with perspective image outputs, ground truth is not required for our training pipeline.

Training datasets: *NYUv2* [53] has a variety of perspective indoor scenes; *VOID* [65] contains indoor perspective office, classroom and stairwell scenes; *IRS* [61] contains rendered perspective scenes of home, restaurant, and store settings; *Hypersim* [52] comprises photorealistic, synthetic images of indoor residential and commercial environments in a perspective reference frame. *Waymo* [56] dataset consists of a diverse set of urban driving scenes.

Test datasets: *ScanNet++* [76] offers 3D reconstructions of diverse indoor scenes, captured using laser scanning and DSLR imaging with a fisheye lens, which allows us to evaluate with real fisheye images and ground truth depth maps. *KITTI-360* [35] includes suburban driving scenes captured with a multi-sensor setup, including fisheye cameras. Notably, it features a different field of view compared to *ScanNet++* [76], allowing us to assess the generalization capability of the Calibration Tokens.

Models. Calibration Tokens do not require specific settings and can be integrated into any model utilizing a vision transformer. We evaluate the effectiveness of our approach using by extending MiDaS [49], DepthAnything [73], and UniDepth [45] to fisheye images. Note that we used 8 tokens per layer for each of the model experiments.

Evaluation Metrics. We evaluate depth prediction accuracy using standard metrics from monocular depth estimation of *root mean squared error* (RMSE) and δ_1 . Details on these metrics can be found in the Supp. Mat.

Implementation Details. We trained our Calibration Tokens based on 3 different FMDEs (MiDaS[49], DepthAnything[49], UniDepth[45]). We utilized the pre-trained ViT-L backbone for MiDaS [49] and DepthAnything [49], and the ViT-S backbone for UniDepth [45]. We trained our model on 4 NVIDIA 3090 GPUs for 40k iterations with a batch size of 16. For input, we used images in the resolution of 518×518 . For testing, we used 462×616 resolution on the *ScanNet++* dataset [76], and 700×700 on the *KITTI-360* dataset [35] to preserve its aspect ratio. We also synthesize random fisheye distortions in the training images. Our Calibration Tokens are trained with a joint dataset consisting of indoor and outdoor datasets totaling up to only 200K samples, and as shown in Tab. 1, obtain comparable results on both domains with fewer samples than existing methods [24] that are trained specifically for each.

4.1. Main Result

We conduct experiments to analyze the impact of Calibration Tokens on model performance. As a baseline, we compare our model to DepthAnyCamera [24], the state-of-the-art monocular depth estimation (MDE) method for fisheye images. Here, we evaluate the DepthAnyCamera model us-

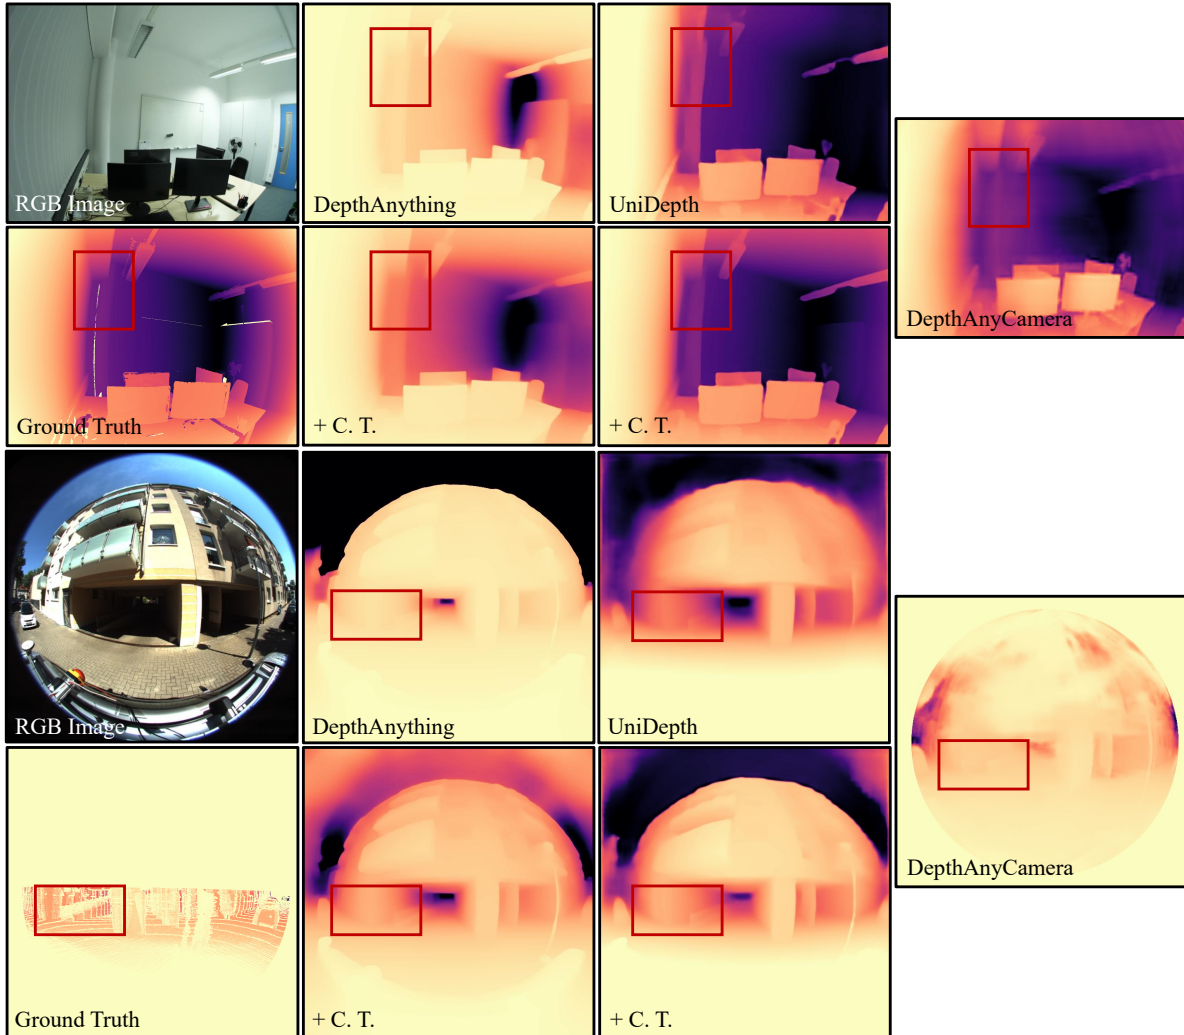


Figure 4. **Comparison on ScanNet++(Indoor) [76] and KITTI-360 [35] dataset.** Qualitative comparison results on ScanNet++ and KITTI-360 datasets. Here, +C. T. indicates prediction results by appending Calibration Tokens to patch embeddings of the model located above. Calibration Tokens enable models to adapt to different fisheye cameras, especially in regions with large distortions.

ing a ResNet101 backbone, and compare both its pretrained model and a model trained on our dataset for any fairness concerns. We also compare with FoVA-Depth [36], which is equirectangular projection based, like DepthAnyCamera. **Indoor Evaluation.** Among the pretrained foundational monocular depth estimators, UniDepth achieves the best performance with our Calibration Tokens on the ScanNet++ indoor dataset as shown in Tab. 1. Notably, our Calibration Tokens enable MiDAS to improve 12% and DepthAnything to achieve a 17% improvement in the RMSE metric compared to the model without Calibration Tokens. Similarly, UniDepth benefits from a 13% improvement in the RMSE metric. Furthermore, compared to the comparison baselines, pretrained DepthAnyCamera and FoVA-Depth, UniDepth with Calibration Tokens surpasses their performance by 11% and 14% in the RMSE metric, respectively. DepthAnyCamera and FoVA-Depth utilize camera intrin-

sics for input images at test time, requiring image transformations back and forth from the equirectangular reference frame, which makes them more error-prone than our direct learning-based approach.

Outdoor Evaluation. We evaluate FMDEs with our proposed Calibration Tokens against state-of-the-art methods in outdoor environments. The results show that Calibration Tokens consistently improve accuracy across different FMDEs in outdoor scenarios. Specifically, MiDAS and DepthAnything achieve improvement in the RMSE metric. UniDepth also improves 2% in the RMSE metric, outperforming the comparison baselines. Given that the KITTI-360 dataset contains highly distorted images with a field of view exceeding 180 degrees, our Calibration Tokens demonstrate robustness across various distortion models.

Our Calibration Tokens are able to outperform DepthAnyCamera and FoVA-Depth without separate

Table 1. **Quantitative comparisons on indoors (ScanNet++) and outdoors (KITTI-360) benchmarks.** We evaluated zero-shot monocular depth estimation by incorporating trained Calibration Tokens into recent foundational monocular depth estimators models. Note: Our method uses the same training set for both the indoor and outdoor settings; whereas existing methods train separate models for each setting.

Testset	Experiment	Model	Train Dataset	RMSE ↓	δ_1 ↑	
ScanNet++ [76]	Baseline	MiDaS [49]	Mix 1.4M	0.506	0.563	
	+ Calibration Tokens	MiDaS [49]	+ Mix 200K	0.446	0.569	
	Baseline	DepthAnything [73]	Mix 63.5M	0.731	0.463	
	+ Calibration Tokens	DepthAnything [73]	+ Mix 200K	0.607	0.506	
	Baseline	UniDepth [45]	Mix 16M	0.279	0.720	
	+ Calibration Tokens	UniDepth [45]	+ Mix 200K	0.244	0.766	
	Comparisons		DepthAnyCamera [24]	Indoor 670K	0.275	0.761
			DepthAnyCamera [24]	Mix 200K	0.761	0.255
			FoVA-Depth [36]	Indoor 190K	0.285	0.548
	KITTI-360 [35]	Baseline	MiDaS [49]	Mix 1.4M	3.312	0.586
+ Calibration Tokens		MiDaS [49]	+ Mix 200K	2.348	0.658	
Baseline		DepthAnything [73]	Mix 63.5M	2.214	0.839	
+ Calibration Tokens		DepthAnything [73]	+ Mix 200K	2.043	0.810	
Baseline		UniDepth [45]	Mix 16M	2.085	0.663	
+ Calibration Tokens		UniDepth [45]	+ Mix 200K	2.040	0.664	
Comparisons			DepthAnyCamera [24]	Outdoor 130K	2.067	0.852
			DepthAnyCamera [24]	Mix 200K	5.675	0.348
			FoVA-Depth [36]	Outdoor 80K	3.096	0.632

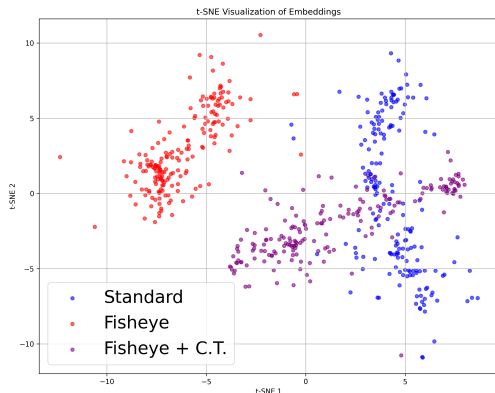


Figure 5. **t-SNE plot of fisheye and perspective embeddings.** Fisheye embeddings become closer to those of perspective images after being modulated by Calibration Tokens.

indoor and outdoor training sets, suggesting the generalization potential of our method to wide ranges of fisheye distortions. Also, the KITTI-360 ground truth points are significantly sparser and more concentrated in ground regions as compared to ScanNet++, which may explain the discrepancy in evaluation metrics. Nonetheless, our method performs comparably without needing to train specialized sets of Calibration Tokens for different fisheye models.

4.2. Analysis

Feature Modulation. To visualize how our Calibration Tokens affect fisheye embeddings, Fig. 5 shows a two-




Figure 6. **Comparison of attention maps for single and multi-layer tokens.** We visualize the attention map of Calibration Tokens to the encoded patch embeddings. Calibration Tokens attend highly to distorted border regions: “Single Token” (top) has little effect in most layers due to lower attention as only a set of tokens are append to the input. The proposed multi-layer or “Layer-wise Tokens” scheme (bottom) attends to embeddings across all layers.

dimensional tSNE reduction to both fisheye and perspective image embeddings from the same set of images. After adding Calibration Tokens to the fisheye embeddings, they are modulated towards the perspective image distribution.

Layer-wise Tokens. As observed in Fig. 6, when we append only a single set of tokens (“Single Token”) at the initial transformer block of the pre-trained model, the Calibration Tokens exhibit limited attention to the patch embeddings across most layers. As a result, the patch embeddings of most layers are unchanged. However, when we attach unique tokens at every layer (“Layer-wise Tokens”), we see

Table 2. **Comparison results with finetuning.** We conducted experiments comparing finetuning (F.T.) with the use of Calibration Tokens (C.T.) added to the baseline model.

Datasets	Models	Exp.	RMSE	δ_1
ScanNet++	MiDAS	F.T.	2.178	0.129
		C.T.	0.446	0.569
	DepthAnything	F.T.	1.459	0.462
		C.T.	0.607	0.506
	UniDepth	F.T.	0.432	0.574
		C.T.	0.244	0.766
KITTI-360	MiDAS	F.T.	9.289	0.098
		C.T.	2.348	0.658
	DepthAnything	F.T.	4.362	0.636
		C.T.	2.043	0.810
	UniDepth	F.T.	3.217	0.403
		C.T.	2.040	0.664

Table 3. **Analysis on computational cost.** We analyze the computational overhead introduced by Calibration Tokens. Values in parentheses indicate the relative increase as a percentage.

Models	Model memory	Tokens memory	Inference time
MiDAS	1.7G	0.8M(0.05%)	0.6ms(0.8%)
DepthAnything	1.7G	0.8M(0.05%)	0.8ms(0.8%)
UniDepth	0.7G	0.2M(0.02%)	0.4ms(0.7%)

higher attention at more layers. Thus, we opt to use the ‘‘Layer-wise’’ approach to better modulate the fisheye patch embeddings toward the distribution of perspective images.

Comparison with Finetuning. To further analyze the robustness of the Calibration Tokens, we conducted experiments comparing our method with a finetuning approach. We trained the model with a fixed learning rate of 10^{-6} on our synthetic fisheye dataset for the same number of iterations. As shown in Tab. 2, the finetuning approach leads to a significant performance drop, highlighting the importance of using Calibration Tokens, which preserve the original model’s training on perspective images.

Computational Cost. Tab. 3 shows the impact of Calibration Tokens on computational costs across different FMDEs. Incorporating Calibration Tokens results in only a 0.05% and 0.02% increase in memory usage, less than 1 MB and a 0.8% and 0.7% increase in inference time, with an added latency of less than 1 ms. This analysis highlights the efficiency of our proposed Calibration Tokens.

4.3. Ablation Study

We conducted an ablation study on different Calibration Token configurations to validate our contributions. Note that ‘‘Single Token’’ refers to a single set of Calibration Tokens

Table 4. **Ablation study.** We ablate the training objective and modulate scheme for our proposed Calibration Tokens.

Dataset	Method	RMSE	δ_1
ScanNet++	Single token	0.260	0.741
	+ LogL1 Loss	0.254	0.752
	+ Layer-wise Tokens	0.244	0.766
KITTI-360	Single token	2.085	0.656
	+ LogL1 Loss	2.065	0.665
	+ Layer-wise Tokens	2.040	0.664

appended in the first layer of the vision transformer without removal, with L1 loss applied. In this configuration, $\phi \in \mathbb{R}^{M \times F}$ as opposed to $\Phi \in \mathbb{R}^{L \times M \times F}$ in the layer-wise setting. The ablation study on the ScanNet++ and KITTI-360 datasets is performed using the UniDepth model.

LogL1 Loss. We observed stable improvements with LogL1 loss compared to baseline L1 loss. As shown in Tab. 4, the LogL1 loss improves both metrics across indoor and outdoor datasets. Qualitative comparisons between L1 and LogL1 objectives are shown in the Supp. Mat.

Layer-wise Tokens. Tab. 4 demonstrates the advantages of using layer-wise tokens over a single set of Calibration Tokens in the first layer. Even when the same number of tokens is fed to the visual transformer layers, we observed a significant improvement in the contribution of layer-wise tokens. This supports our hypothesis about how the influence of Calibration Tokens diminishes through a forward pass as observed in our experiments by appending a single set of tokens at the first layer. Fig. 6 visualizes attention.

5. Discussion

Calibration Tokens enable FMEs to adapt to images captured by fisheye cameras. Empirically, our method improves on monocular depth estimation on fisheye cameras. While our method trains only one set of tokens for both indoor and outdoor settings, our promising results motivates this as a general approach to adapting vision foundational models. Furthermore, a convenience afforded by our method is in the reuse and backward-compatibility of FMDEs with perspective images. This reduces the operational overhead of multi-camera systems by enabling a single FMDE to handle multiple camera inputs – adding cameras become as easy as appending tokens.

Limitations. While we offer a light-weight method of extending FMDEs to fisheye images, its success inherently depends on the quality and representational power of the underlying FMDEs. If the pretrained model struggles with certain 3D scenes or lighting conditions for perspective images, then these issues carry over. Nonetheless, as novel FMDEs emerge, our framework can be readily transferred to new models using transformer-based architectures.

Acknowledgments This work is supported by NSF 2112562 Athena AI Institute and the Rosenfeld Science Scholars Program.

References

- [1] Bruno Arsenali, Prashanth Viswanath, and Jelena Novosel. Rotinvm1: Rotation invariant multinet on fisheye images for autonomous driving applications. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. 3
- [3] Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. In *International Conference on Learning Representations*, 2025. 3
- [4] Mikhail S Burtsev, Yuri Kuratov, Anton Peganov, and Grigory V Sapunov. Memory transformer. *arXiv preprint arXiv:2006.11527*, 2020. 4
- [5] Wenjie Chang, Yueyi Zhang, and Zhiwei Xiong. Transformer-based monocular depth estimation with attention supervision. In *32nd British Machine Vision Conference (BMVC 2021)*, 2021. 3
- [6] Xien Chen, Suchisrit Gangopadhyay, Michael Chu, Patrick Rim, Hyungseob Park, and Alex Wong. Uncle: Unsupervised continual learning of depth completion. *arXiv preprint arXiv:2410.18074*, 2024. 3
- [7] Hyesong Choi, Hunsang Lee, Sunkyung Kim, Sunok Kim, Seungryong Kim, Kwanghoon Sohn, and Dongbo Min. Adaptive confidence thresholding for monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12808–12818, 2021. 3
- [8] Younjoon Chung, Hyungseob Park, Patrick Rim, Xiaoran Zhang, Jihe He, Ziyao Zeng, Safa Cicek, Byung-Woo Hong, James S. Duncan, and Alex Wong. Eta: Energy-based test-time adaptation for depth completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. 3
- [9] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023. 4
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. 3
- [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*, 2021. 2, 3
- [12] C Brown Duane. Close-range camera calibration. *Photogramm. Eng.*, 37(8):855–866, 1971. 3
- [13] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*, pages 2650–2658, 2015. 3
- [14] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014. 3
- [15] Vadim Ezhov, Hyungseob Park, Zhaoyang Zhang, Rishi Upadhyay, Howard Zhang, Chethan Chinder Chandrappa, Achuta Kadambi, Yunhao Ba, Julie Dorsey, and Alex Wong. All-day depth completion. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024. 3
- [16] Xiaohan Fei, Alex Wong, and Stefano Soatto. Geosupervised visual depth prediction. *IEEE Robotics and Automation Letters*, 4(2):1661–1668, 2019. 3
- [17] Hao Feng, Wendi Wang, Jiajun Deng, Wengang Zhou, Li Li, and Houqiang Li. Simfir: A simple framework for fisheye image rectification with self-supervised representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12418–12427, 2023. 3, 4
- [18] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2002–2011, 2018. 3
- [19] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 740–756. Springer, 2016. 3
- [20] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 270–279, 2017. 3
- [21] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3828–3838, 2019. 3
- [22] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. In *International Conference on Learning Representations*, 2019. 3
- [23] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Rantos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2485–2494, 2020. 3
- [24] Yuliang Guo, Sparsh Garg, S Mahdi H Miangoleh, Xinyu Huang, and Liu Ren. Depth any camera: Zero-shot met-

- ric depth estimation from any camera. *arXiv preprint arXiv:2501.02464*, 2025. 3, 5, 7
- [25] Juho Kannala and Sami Brandt. A generic camera calibration method for fish-eye lenses. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, pages 10–13. IEEE, 2004. 3, 4
- [26] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9492–9502, 2024. 3
- [27] R Kingslake. A history of the photographic lens. *University of Rochester NY*, 145, 1989. 3
- [28] Varun Ravi Kumar, Marvin Klingner, Senthil Yogamani, Stefan Milz, Tim Fingscheidt, and Patrick Mader. Syndistnet: Self-supervised monocular fisheye camera distance estimation synergized with semantic segmentation for autonomous driving. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 61–71, 2021. 3
- [29] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *2016 Fourth international conference on 3D vision (3DV)*, pages 239–248. IEEE, 2016. 3
- [30] Dong Lao, Yangchao Wu, Tian Yu Liu, Alex Wong, and Stefano Soatto. Sub-token vit embedding via stochastic resonance transformers. In *International Conference on Machine Learning*. PMLR, 2024.
- [31] Dong Lao, Fengyu Yang, Daniel Wang, Hyoungseob Park, Samuel Lu, Alex Wong, and Stefano Soatto. On the viability of monocular depth pre-training for semantic segmentation. In *European Conference on Computer Vision*. Springer, 2024.
- [32] Bo Li, Chunhua Shen, Yuchao Dai, Anton Van Den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1119–1127, 2015. 3
- [33] Zhenyu Li, Zehui Chen, Xianming Liu, and Junjun Jiang. Depthformer: Exploiting long-range correlation and local information for accurate monocular depth estimation. *Machine Intelligence Research*, 20(6):837–854, 2023. 3
- [34] Kang Liao, Chunyu Lin, Yao Zhao, and Moncef Gabbouj. Dr-gan: Automatic radial distortion rectification using conditional gan in real-time. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(3):725–733, 2019. 3
- [35] Yiyi Liao, Jun Xie, and Andreas Geiger. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3292–3310, 2022. 5, 6, 7
- [36] Daniel Lichy, Hang Su, Abhishek Badki, Jan Kautz, and Orazio Gallo. Fova-depth: Field-of-view agnostic depth estimation for cross-dataset generalization. In *2024 International Conference on 3D Vision (3DV)*, pages 1–10. IEEE, 2024. 3, 6, 7
- [37] 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*, 38(10):2024–2039, 2015. 3
- [38] Tian Yu Liu, Parth Agrawal, Allison Chen, Byung-Woo Hong, and Alex Wong. Monitored distillation for positive congruent depth completion. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 35–53. Springer, 2022. 3
- [39] Xiaoyang Lyu, Liang Liu, Mengmeng Wang, Xin Kong, Lina Liu, Yong Liu, Xinxin Chen, and Yi Yuan. Hr-depth: High resolution self-supervised monocular depth estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2294–2301, 2021. 3
- [40] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5667–5675, 2018. 3
- [41] Kenro Miyamoto. Fish eye lens. *JOSA*, 54(8):1060–1061, 1964. 3
- [42] Hyoungseob Park, Anjali Gupta, and Alex Wong. Test-time adaptation for depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20519–20529, 2024. 3
- [43] Rui Peng, Ronggang Wang, Yawen Lai, Luyang Tang, and Yangang Cai. Excavating the potential capacity of self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15560–15569, 2021. 5
- [44] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10106–10116, 2024. 1, 2, 3
- [45] Luigi Piccinelli, Christos Sakaridis, Yung-Hsu Yang, Mattia Segu, Siyuan Li, Wim Abbeels, and Luc Van Gool. Unidepthv2: Universal monocular metric depth estimation made simpler. *arXiv preprint arXiv:2502.20110*, 2025. 5, 7
- [46] Matteo Poggi, Fabio Tosi, and Stefano Mattoccia. Learning monocular depth estimation with unsupervised trinocular assumptions. In *2018 International conference on 3d vision (3DV)*, pages 324–333. IEEE, 2018. 3
- [47] Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia. On the uncertainty of self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3227–3237, 2020. 3
- [48] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 3
- [49] Rene 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 & Machine Intelligence*, 44(03):1623–1637, 2022. 2, 3, 5, 7
- [50] Patrick Rim, Hyoungeob Park, Vadim Ezhov, Jeffrey Moon, and Alex Wong. Radar-guided polynomial fitting for metric depth estimation. *arXiv preprint arXiv:2503.17182*, 2025. 3
- [51] Patrick Rim, Hyoungeob Park, Ziyao Zeng, Younjoon Chung, and Alex Wong. Protodepth: Unsupervised continual depth completion with prototypes. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6304–6316, 2025. 3
- [52] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021. 5
- [53] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *European Conference on Computer Vision*, 2012. 5
- [54] Akash Deep Singh, Yunhao Ba, Ankur Sarker, Howard Zhang, Achuta Kadambi, Stefano Soatto, Mani Srivastava, and Alex Wong. Depth estimation from camera image and mmwave radar point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9275–9285, 2023. 3
- [55] Daniel E Stevenson and Margaret M Fleck. Nonparametric correction of distortion. In *Proceedings Third IEEE Workshop on Applications of Computer Vision. WACV'96*, pages 214–219. IEEE, 1996. 3
- [56] Pei Sun, Henrik Kretschmar, 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*, pages 2446–2454, 2020. 5
- [57] Fabio Tosi, Filippo Aleotti, Matteo Poggi, and Stefano Mattoccia. Learning monocular depth estimation infusing traditional stereo knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9799–9809, 2019. 3
- [58] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers and distillation through attention. In *Proceedings of the 38th International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 3
- [59] Rishi Upadhyay, Howard Zhang, Yunhao Ba, Ethan Yang, Blake Gella, Sicheng Jiang, Alex Wong, and Achuta Kadambi. Enhancing diffusion models with 3d perspective geometry constraints. *ACM Transactions on Graphics (TOG)*, 42(6):1–15, 2023. 3
- [60] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2022–2030, 2018. 3
- [61] Qiang Wang, Shizhen Zheng, Qingsong Yan, Fei Deng, Kaiyong Zhao, and Xiaowen Chu. Irs: A large naturalistic indoor robotics stereo dataset to train deep models for disparity and surface normal estimation. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2021. 5
- [62] Jamie Watson, Michael Firman, Gabriel J Brostow, and Daniyar Turmukhambetov. Self-supervised monocular depth hints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2162–2171, 2019. 3
- [63] Alex Wong and Stefano Soatto. Bilateral cyclic constraint and adaptive regularization for unsupervised monocular depth prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5644–5653, 2019.
- [64] Alex Wong, Safa Cicek, and Stefano Soatto. Targeted adversarial perturbations for monocular depth prediction. *Advances in neural information processing systems*, 33:8486–8497, 2020.
- [65] Alex Wong, Xiaohan Fei, Stephanie Tsuei, and Stefano Soatto. Unsupervised depth completion from visual inertial odometry. *IEEE Robotics and Automation Letters*, 5(2): 1899–1906, 2020. 5
- [66] Alex Wong, Safa Cicek, and Stefano Soatto. Learning topology from synthetic data for unsupervised depth completion. *IEEE Robotics and Automation Letters*, 6(2):1495–1502, 2021.
- [67] Alex Wong, Xiaohan Fei, Byung-Woo Hong, and Stefano Soatto. An adaptive framework for learning unsupervised depth completion. *IEEE Robotics and Automation Letters*, 6(2):3120–3127, 2021. 3
- [68] Yangchao Wu, Tian Yu Liu, Hyoungeob Park, Stefano Soatto, Dong Lao, and Alex Wong. Augundo: Scaling up augmentations for monocular depth completion and estimation. In *European Conference on Computer Vision*, pages 274–293. Springer, 2024. 3, 5
- [69] Chao Xia, Chenfeng Xu, Patrick Rim, Mingyu Ding, Nanning Zheng, Kurt Keutzer, Masayoshi Tomizuka, and Wei Zhan. Quadric representations for lidar odometry, mapping and localization. *IEEE Robotics and Automation Letters*, 8(8):5023–5030, 2023. 3
- [70] Yichen Xie, Chenfeng Xu, Marie-Julie Rakotosaona, Patrick Rim, Federico Tombari, Kurt Keutzer, Masayoshi Tomizuka, and Wei Zhan. Sparsefusion: Fusing multi-modal sparse representations for multi-sensor 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17591–17602, 2023. 3
- [71] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5354–5362, 2017. 3
- [72] Fengyu Yang, Chao Feng, Ziyang Chen, Hyoungeob Park, Daniel Wang, Yiming Dou, Ziyao Zeng, Xien Chen, Rit Gangopadhyay, Andrew Owens, and Alex Wong. Binding touch to everything: Learning unified multimodal tactile representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26340–26353, 2024. 3

- [73] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2, 3, 5, 7
- [74] Yanchao Yang, Alex Wong, and Stefano Soatto. Dense depth posterior (ddp) from single image and sparse range. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3353–3362, 2019. 3
- [75] Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, and Ram Nevatia. Lego: Learning edge with geometry all at once by watching videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 225–234, 2018. 3
- [76] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023. 5, 6, 7
- [77] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5684–5693, 2019. 3
- [78] Xiaoqing Yin, Xinchao Wang, Jun Yu, Maojun Zhang, Pascal Fua, and Dacheng Tao. Fisheyecnet: A multi-context collaborative deep network for fisheye image rectification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 469–484, 2018. 4
- [79] Senthil Yogamani, David Unger, Venkatraman Narayanan, and Varun Ravi Kumar. Fisheyebevseg: Surround view fisheye cameras based bird’s-eye view segmentation for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1331–1334, 2024. 3
- [80] 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*, pages 3916–3925, 2022. 3
- [81] Ziyao Zeng, Jingcheng Ni, Daniel Wang, Patrick Rim, Younjoon Chung, Fengyu Yang, Byung-Woo Hong, and Alex Wong. Priordiffusion: Leverage language prior in diffusion models for monocular depth estimation. *arXiv preprint arXiv:2411.16750*, 2024. 3
- [82] Ziyao Zeng, Daniel Wang, Fengyu Yang, Hyungseob Park, Stefano Soatto, Dong Lao, and Alex Wong. Worddepth: Variational language prior for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9708–9719, 2024.
- [83] Ziyao Zeng, Yangchao Wu, Hyungseob Park, Daniel Wang, Fengyu Yang, Stefano Soatto, Dong Lao, Byung-Woo Hong, and Alex Wong. Rsa: Resolving scale ambiguities in monocular depth estimators through language descriptions. *Advances in neural information processing systems*, 37, 2024. 3
- [84] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 340–349, 2018. 3
- [85] Ning Zhang, Francesco Nex, George Vosselman, and Norman Kerle. Lite-mono: A lightweight cnn and transformer architecture for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18537–18546, 2023. 3
- [86] Chaoqiang Zhao, Youmin Zhang, Matteo Poggi, Fabio Tosi, Xianda Guo, Zheng Zhu, Guan Huang, Yang Tang, and Stefano Mattoccia. Monovit: Self-supervised monocular depth estimation with a vision transformer. In *2022 International Conference on 3D Vision (3DV)*, pages 668–678. IEEE, 2022. 3
- [87] Guoyang Zhao, Yuxuan Liu, Weiqing Qi, Fulong Ma, Ming Liu, and Jun Ma. Fisheyedepth: A real scale self-supervised depth estimation model for fisheye camera. *arXiv preprint arXiv:2409.15054*, 2024. 3
- [88] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1851–1858, 2017. 3