

# VGGT-360: Geometry-Consistent Zero-Shot Panoramic Depth Estimation

Jiayi Yuan<sup>1</sup>, Haobo Jiang<sup>2</sup>, De Wen Soh<sup>1</sup>, Na Zhao<sup>\*1</sup>

<sup>1</sup> Singapore University of Technology and Design, <sup>2</sup> Nanyang Technological University

jiayi-yuan@mymail.sutd.edu.sg, haobo.jiang@ntu.edu.sg, dewen\_soh@sutd.edu.sg,  
na\_zhao@sutd.edu.sg

## Abstract

This paper presents VGGT-360, a novel training-free framework for zero-shot, geometry-consistent panoramic depth estimation. Unlike prior view-independent training-free approaches, VGGT-360 reformulates the task as panoramic reprojection over multi-view reconstructed 3D models by leveraging the intrinsic 3D consistency of VGGT-like foundation models, thereby unifying fragmented per-view reasoning into a coherent panoramic understanding. To achieve robust and accurate estimation, VGGT-360 integrates three plug-and-play modules that form a unified panorama-to-3D-to-depth framework: (i) Uncertainty-guided adaptive projection slices panoramas into perspective views to bridge the domain gap between panoramic inputs and VGGT’s perspective prior. It estimates gradient-based uncertainty to allocate denser views to geometry-poor regions, yielding geometry-informative inputs for VGGT. (ii) Structure-saliency enhanced attention strengthens VGGT’s robustness during 3D reconstruction by injecting structure-aware confidence into its attention layers, guiding focus toward geometrically reliable regions and enhancing cross-view coherence. (iii) Correlation-weighted 3D model correction refines the reconstructed 3D model by reweighting overlapping points using attention-inferred correlation scores, providing a consistent geometric basis for accurate panoramic reprojection. Extensive experiments show that VGGT-360 outperforms both trained and training-free state-of-the-art methods across multiple resolutions and diverse indoor and outdoor datasets. The code is available at <https://github.com/Yuanjiayi/VGGT-360>.

## 1. Introduction

Monocular depth estimation (MDE) from 360° panoramic images is an emerging research problem that plays a criti-

\* Corresponding author.

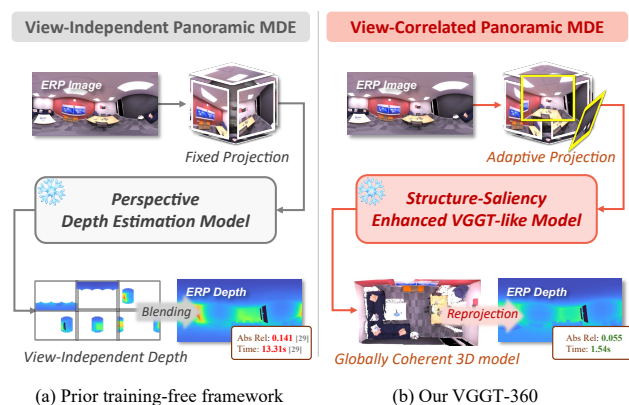


Figure 1. Comparison between the conventional training-free panoramic depth estimation framework and our VGGT-360. Unlike view-independent inference methods (e.g., 360MD [27]), VGGT-360 reconstructs a globally coherent 3D representation via VGGT-like 3D foundation models and reprojects it to the panorama, unifying fragmented per-view predictions into consistent, cross-view correlated depth with superior performance.

cal role in omnidirectional perception. Unlike conventional perspective depth, panoramic depth provides complete geometric context, which is essential for applications such as omnidirectional SLAM [10, 47], virtual reality [4, 18], and autonomous navigation systems [11]. However, panoramic MDE faces two major challenges: **i)** Panoramas are typically represented by the equirectangular projection (ERP), which flattens a sphere into a 2D image and inevitably introduces severe geometric distortions, thereby making conventional perspective MDE models unsuitable for direct application [12]; and **ii)** The acquisition of large-scale annotated panoramic datasets is extremely difficult and costly, which severely limits the performance and generalization capability of existing supervised approaches [27, 39].

Existing panoramic MDE methods can be broadly categorized into *training-based* methods and *training-free* methods. Training-based approaches [7, 41, 44, 46] pre-

dict depth directly from ERP panoramas using sphere- or distortion-aware MDE models. However, their performance is limited by the scarcity of labeled panoramic data, resulting in lower accuracy and poor generalization. To address this, training-free methods [20, 24, 27] decompose the panorama into perspective views, infer depth independently via pre-trained perspective MDE models, and fuse them into the full ERP map (Fig. 1(a)). Nonetheless, this strategy suffers from geometric inconsistency, as view-independent inference lacks cross-view interaction, leading to scale ambiguity and depth discontinuities across views, which in turn degrade geometric fidelity and structural detail.

In this paper, beyond the conventional view-independent, training-free paradigm for panoramic MDE, we reformulate the task as panoramic reprojection over multi-view reconstructed, globally consistent 3D models, as illustrated in Fig. 1(b). Inspired by the success of VGGT [38] in capturing intrinsic global geometry across multiple views, we introduce **VGGT-360**, a training-free, geometry-consistent panoramic MDE framework that leverages the holistic 3D reasoning capabilities of VGGT-like 3D foundation models to unify fragmented per-view reasoning into a coherent panoramic understanding. However, since VGGT-like models (e.g., VGGT [38],  $\pi^3$  [40]) are originally trained on perspective images, directly extending them to panoramic scenes introduces a severe domain gap, making effective adaptation non-trivial. The network must simultaneously handle spherical distortion, non-uniform resolution, and 360° wrap-around continuity, while producing geometrically consistent depth under unseen panorama scenarios.

To address these limitations, we innovatively design VGGT-360 as a three-stage, training-free framework comprising **i)** uncertainty-guided adaptive projection, **ii)** structure-saliency enhanced attention, and **iii)** correlation-weighted 3D correction, jointly improving geometric fidelity and reconstruction robustness. Specifically, *uncertainty-guided adaptive projection* introduces gradient-derived geometric uncertainty to adaptively slice the panorama into perspective views. Unlike fixed projection schemes [20, 27], our adaptive method allocates denser sampling to geometry-deficient regions under coverage-and-overlap constraints, yielding more reliable, geometry-informative multi-view inputs for VGGT. Furthermore, despite the remarkable zero-shot generalization of VGGT-like foundation models, their performance often degrades in weakly structured areas, compromising reconstruction quality. To mitigate this, *structure-saliency enhanced attention* injects a structure-saliency confidence map into VGGT’s attention layers, guiding the model toward reliable geometric and boundary cues for stable 3D reasoning. Finally, *correlation-weighted 3D model correction* refines the reconstructed model by assigning correlation-based reliability scores to overlapping 3D points, reinforcing consistent

structures while suppressing ambiguous ones.

Notably, our VGGT-360 is highly general and plug-and-play, enabling seamless integration with diverse VGGT variants such as  $\pi^3$  [40] and Fastvggt [29] for enhanced performance. Benefiting from the strong generalization of VGGT-like foundation models, VGGT-360 also exhibits remarkable zero-shot capability, delivering reliable depth estimation across varying resolutions and both indoor and outdoor scenes without any fine-tuning.

Our main contributions are summarized as follows:

- We propose **VGGT-360**, a novel training-free, 3D model-aware framework that exploits the global consistency of VGGT-like 3D foundation models for coherent panoramic depth estimation, significantly outperforming prior view-independent panoramic MDE approaches.
- We effectively adapt VGGT-like models for panorama-driven 3D reconstruction via three novel plug-and-play modules, *i.e.*, uncertainty-guided adaptive projection, structure-saliency enhanced attention, and correlation-weighted 3D model correction, together enabling robust and structurally consistent panoramic depth prediction.
- Extensive experiments across multiple resolutions and diverse indoor/outdoor datasets validate the effectiveness and zero-shot generalization ability of VGGT-360. Notably, it surpasses previous SOTA methods by 27–36% in Abs Rel on Stanford2D3D [3] and Replica360-2K [30].

## 2. Related Work

**Panoramic Depth Estimation.** Monocular depth estimation (MDE) for 360° panoramic images aims to recover global scene depth from equirectangular projection (ERP) images. Existing learning-based methods mainly follow two directions. One designs sphere-aware networks that operate directly in the ERP domain to handle latitude-dependent distortions via distortion-aware convolutions or spherical positional encodings [7, 12, 31, 33, 44, 46, 49]. The other adopts multi-projection fusion frameworks, which decompose panoramas into multiple perspective views to complement ERP processing and alleviate projection distortions [2, 19, 21, 23, 35–37]. However, both categories rely on supervised learning and are fundamentally limited by the scarcity of annotated panoramic datasets, leading to weak generalization in diverse scenes [20, 39].

**Zero-Shot Monocular Depth Estimation.** Recent monocular depth foundation models, such as MiDaS [26], Depth Anything [42], Metric3D [43], and Omnidata [13], have demonstrated strong zero-shot generalization to unseen domains and scenes without fine-tuning. Inspired by this, recent panoramic MDE approaches [20, 24, 27, 39] have adopted these models as powerful pre-trained depth priors. For instance, 360MD [27] first applied pre-trained perspective models to panoramic inputs by slicing the panorama into multiple views and blending their predictions through

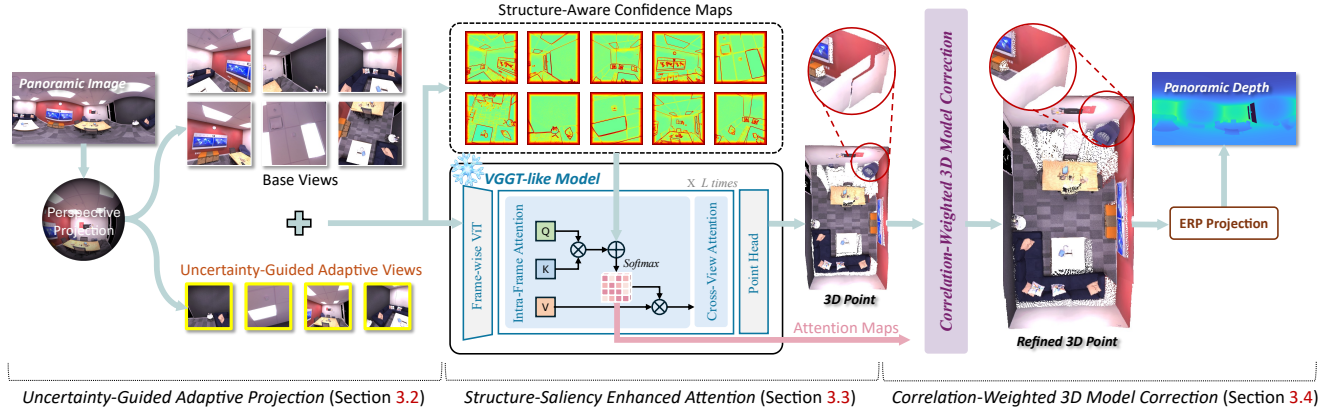


Figure 2. **Framework Overview of VGGT-360.** Given a panoramic image, we first perform *uncertainty-guided adaptive projection* to produce geometry-informative views for VGGT. With *structure-saliency enhanced attention*, VGGT reconstructs a structure-faithful 3D model, which is then refined by *correlation-weighted 3D model correction* and reprojected into a globally consistent panoramic depth map.

optimization-based fusion. Depth Anywhere [39] and RPG360 [20] further introduced distillation and geometric optimization to enhance panoramic integration. However, these zero-shot frameworks typically infer each view independently without enforcing cross-view geometric consistency, resulting in scale discrepancies and fragmented 3D structures. This limitation motivates a training-free yet geometry-consistent framework that transforms fragmented per-view predictions into a unified panoramic depth map.

**Visual Geometry Grounded Transformer.** VGGT [38] is a unified architecture for reconstructing 3D models from multi-view images. Unlike conventional depth foundation models [6, 8, 13, 42], which learn implicit geometric priors from large-scale 2D datasets, VGGT explicitly reconstructs consistent 3D representations by leveraging cross-view geometric cues. Subsequent VGGT variants [29, 34, 40] have further enhanced its efficiency and robustness. For example, Fastvggt [29] accelerates inference with lightweight attention, while  $\pi^3$  [40] eliminates reference-view dependency to enable permutation-equivariant reconstruction. Other extensions adopt VGGT as a geometry-aware prior for dense novel view synthesis [22] and robotic perception [14]. In this work, we are the first to extend VGGT-like foundation models to panoramic depth, marking a significant step toward training-free omnidirectional 3D reasoning.

### 3. Method

In this paper, we present *VGGT-360*, a novel training-free framework for geometry-consistent panoramic depth estimation. As illustrated in Fig. 2, given a single equirectangular panorama  $\mathcal{I}_{\text{erp}} \in \mathbb{R}^{H \times W \times 3}$ , VGGT-360 adaptively projects it into multiple perspective views, performs multi-view 3D reasoning through VGGT-like 3D foundation models (e.g., VGGT [38],  $\pi^3$  [40], and Fastvggt [29]) to reconstruct a globally coherent 3D representation, and finally reprojects it back to the panorama to produce the depth map

$\mathcal{D}_{\text{erp}} \in \mathbb{R}^{H \times W}$  via ERP projection.

To realize this process effectively, VGGT-360 comprises three key training-free modules that enhance the performance of the VGGT-like model for panoramic input: **i) Uncertainty-guided adaptive projection**, which generates perspective images for VGGT by allocating denser sampling to geometrically uncertain regions while preserving global spherical continuity (Section 3.2); **ii) Structure-saliency enhanced attention**, which enhances VGGT’s attention by emphasizing structurally consistent geometric features, thereby improving 3D reasoning and suppressing artifacts in geometrically ambiguous regions (Section 3.3); and **iii) Correlation-weighted 3D model correction**, which refines the 3D model by reweighting points in overlapping regions based on correlation cues derived from VGGT’s inter-frame attention, ensuring high-quality panoramic depth estimation (Section 3.4).

#### 3.1. VGGT-like 3D Foundation Models

To enable multi-view reasoning, our framework leverages VGGT-like 3D Foundation Models as a core component. Given multi-view perspective inputs  $\mathcal{V}_{\text{per}}$ , VGGT-like models tokenize images into patch embeddings and apply alternating intra-frame and cross-view attention for 3D reasoning, directly predicting camera parameters, point maps, depth, and point tracks. We adopt the point head output and reproject it into panoramic depth via ERP projection.

#### 3.2. Uncertainty-Guided Adaptive Projection

To provide the perspective multi-view inputs required by VGGT, we first slice the equirectangular panoramic image  $\mathcal{I}_{\text{erp}}$  into a set of perspective views  $v \in \mathcal{V}_{\text{per}}$ . Prior methods [20, 27, 36] typically adopt uniform, pre-defined projection schemes (e.g., cubemap), under the assumption that all viewing directions are treated equally, regardless of their geometric informativeness. In practice, views dom-

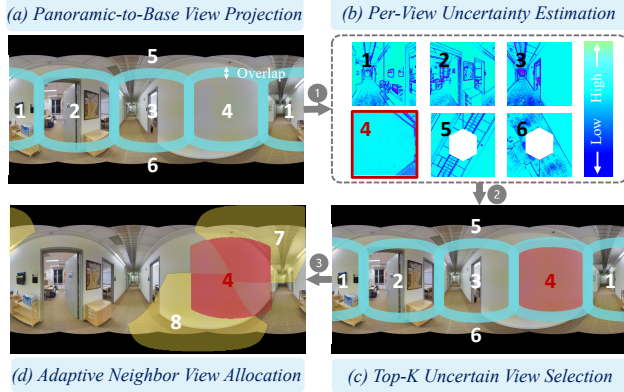


Figure 3. **Pipeline of our uncertainty-guided adaptive projection.** We first generate  $N_B$  base views from the panorama, compute per-view uncertainty maps via edge-based scoring, and select the top- $K$  most uncertain views (with  $N_B=6$ ,  $K=1$  in this example). These views are then augmented with neighboring projections to form a geometry-aware multi-view set as input for VGGT.

inated by weakly structured regions (*e.g.*, planar walls or ceilings) offer limited geometric cues, making depth reasoning more challenging. To address this, we propose *uncertainty-guided adaptive projection*, which dynamically allocates denser views to regions with high geometric ambiguity, enabling VGGT to enrich its geometric understanding via neighboring views. Our projection strategy consists of two stages: *uncertainty-guided scoring*, which quantifies geometric ambiguity across base views, and *adaptive neighbor augmentation*, which allocates additional view samples based on these uncertainties.

**Uncertainty-Guided Scoring.** As shown in Fig. 3(a-b), we begin with a pre-defined perspective projection using  $N_B$  ( $N_B \geq 6$ ) base views  $\mathcal{B} = \{v_b^n\}_{n=1}^{N_B}$ , ensuring full coverage with controlled overlap. To identify geometry-sparse views for subsequent denser allocation, we assign each view an uncertainty score reflecting its geometric informativeness. Given that gradient magnitude effectively captures edge richness and geometric saliency, we compute a per-pixel *uncertainty map*  $\mathbf{U}(p) = \sigma(-\mathbf{Z}(p))$  from the grayscale image of each base view using the Sobel operator:

$$\mathbf{Z}(p) = (\mathbf{G}(p) - \text{median}_{p' \in \Omega(v_b)}(\mathbf{G}(p'))) / \tau, \quad (1)$$

where  $\mathbf{G}(p)$  denotes the Sobel gradient magnitude at pixel  $p$ ,  $\tau$  controls normalization sensitivity, and  $\Omega(v_b)$  represents the valid region of base view  $v_b$ . The area-weighted uncertainty score of view  $v_b$  is computed as:

$$\mathbf{S}(v_b) = \frac{\sum_{p \in \Omega(v_b)} \mathbf{1}_{\text{valid}}(p) \mathbf{U}(p)}{\sum_{p \in \Omega(v_b)} \mathbf{1}_{\text{valid}}(p)}, \quad (2)$$

where  $\mathbf{1}_{\text{valid}}(p)$  is an indicator function that masks out invalid pixels in perspective views.

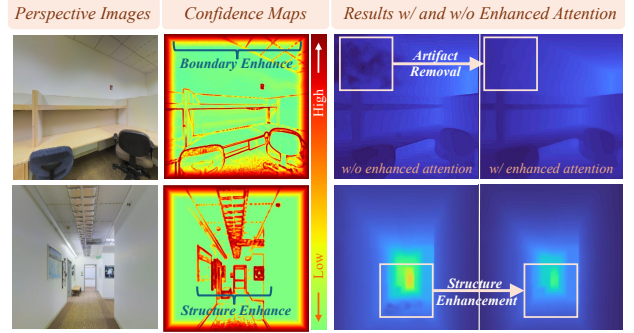


Figure 4. **Comparison of results before and after applying our structure-saliency enhanced attention mechanism.** Guided by our well-designed *structure-aware confidence map*, our VGGT-360 effectively removes artifacts and preserves geometric structures in weakly structured regions, which are easily affected by illumination cues and noise.

**Adaptive Neighbor Augmentation.** As shown in Fig. 3(c-d), given the set of uncertainty scores  $\{\mathbf{S}(v_b)\}_{v_b \in \mathcal{B}}$ , we perform adaptive neighbor-augmented sampling by selecting the top- $K$  base views with the highest uncertainty:  $\mathcal{B}^* = \text{Top-}K(\{\mathbf{S}(v_b)\}_{v_b \in \mathcal{B}})$ . For each selected view  $v_b^* \in \mathcal{B}^*$ , we generate two neighboring views  $\mathcal{N}(v_b^*)$  with predefined yaw and pitch offsets (*i.e.*, upper-right and lower-left of the view center). This adaptive strategy yields an efficient perspective view set:  $\mathcal{V}_{\text{per}} = \mathcal{B} \cup \mathcal{N}(\mathcal{B}^*)$ , which maintains global coverage while adaptively increasing sampling density in geometrically ambiguous regions under a limited budget.

### 3.3. Structure-Saliency Enhanced Attention

Once the perspective view set is obtained, we leverage the powerful 3D-consistent reconstruction capability of VGGT-like foundation models to aggregate multiple perspective views into a globally coherent 3D representation, serving as the foundation for panoramic MDE. However, we observe that these models often deteriorate in weakly structured regions lacking reliable geometric cues, leading to artifacts and hallucinatory depth predictions (Fig. 4), despite their strong zero-shot generalization ability. To address this issue, we introduce a *structure-saliency enhanced attention* mechanism that integrates structure-derived priors (*i.e.*, structure-aware confidence maps) into the intra-frame attention layers of VGGT-like models, steering multi-view aggregation toward geometrically stable and reliable regions.

**Structure-Aware Confidence Map.** To inject structural priors into VGGT’s attention without altering its pretrained weights, we define a pixel-level *structure-aware confidence map*  $\mathbf{M}_s$  for each view. As shown in Fig. 4,  $\mathbf{M}_s$  integrates structural saliency and view boundary-aware weighting to strengthen attention in geometry-reliable regions while preserving feature continuity across overlaps, thereby improving both local depth fidelity and global coherence.

Specifically, we first compute the gradient-based ge-

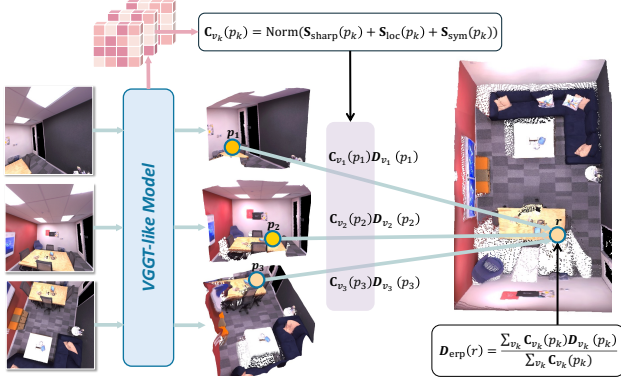


Figure 5. **Pipeline of our correlation-weighted 3D model correction module.** Each overlapping 3D point is assigned a correlation score derived from VGGT’s intra-frame attention. These scores are then used as reliability weights to refine the reconstructed 3D model and produce the final coherent ERP depth.

ometric prior  $\mathbf{M}_g$  using the Sobel operator  $\mathbf{Z}(p)$  (see Eq.1) on the grayscale image to highlight structurally reliable regions:  $\mathbf{M}_g(p) = \sigma(\mathbf{Z}(p))$ . To encourage stronger cross-view interaction near overlapping boundaries, we introduce an *edge-band prior* defined as  $\mathbf{E}(p) = 1 [\max(|x|, |y|) \geq 1 - m]$ , emphasizing uncertain pixels near the view edges. Here,  $(x, y) \in [-1, 1]^2$  are the normalized image coordinates of pixel  $p$ , and  $m$  controls the edge-band width. The final confidence map is computed as:

$$\mathbf{M}_s(p) = \mathbf{1}_{\text{valid}}(p) \cdot [(1 - \mathbf{E}(p)) \cdot \mathbf{M}_g(p) + \mathbf{E}(p)]. \quad (3)$$


**Structure-Saliency Enhanced Frame Attention.** Given the proposed confidence prior  $\mathbf{M}_s$ , we enhance intra-frame attention of VGGT-like models to emphasize geometrically reliable keys while suppressing ambiguous ones. Let  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  denote the query, key, and value features, respectively. We incorporate the confidence map  $\mathbf{M}_s$  as an additive log-confidence bias to the attention scores:

$$\mathbf{M}_{\text{Attn}} = \text{softmax}\left(\mathbf{Q}\mathbf{K}^\top / \sqrt{d} + \log(\mathbf{M}_s)\right). \quad (4)$$

As shown in Fig. 4, this bias explicitly steers attention toward structurally stable regions, while effectively suppressing spurious artifacts in geometrically ambiguous areas, without altering the pretrained backbone weights.

### 3.4. Correlation-Weighted 3D Model Correction

Building on the reconstructed 3D representation in Sec. 3.3, we introduce a *correlation-weighted 3D model correction* module to further enforce geometric consistency across overlapping regions, as shown in Fig. 5. Our core idea is to assign each overlapping 3D point a correlation score, derived from intra-frame attention, that reflects its connectivity with neighboring points. Points with higher correlations are regarded as more reliable and thus receive greater weights for the final ERP depth construction.



(a) 3D Model w/o our correction module (b) 3D Model with our correction module

Figure 6. **Comparison of reconstructed 3D models without (a) and with (b) our correlation-weighted 3D model correction.** Our correction module significantly enhances surface continuity and removes artifacts in overlapping regions.

Specifically, consider an ERP pixel  $r$  that is observed by a set of  $N_K$  perspective views, denoted as  $\mathcal{V}_{\text{per}}^k = \{v_1, \dots, v_{N_K}\}$ . For each view  $v_k$ , we project the ERP pixel  $r$  onto its corresponding 3D point  $p_k$  to obtain the depth observation  $D_{v_k}(p_k)$ . Subsequently, a correlation weight  $C_{v_k}(p_k)$  for the point  $p_k$  is derived from that view’s attention map  $\widetilde{\mathbf{M}}_{\text{Attn}}$  of the final intra-frame attention layer. Finally, the final ERP depth  $D_{\text{erp}}$  is computed via a pixel-wise weighted aggregation of these multi-view depths:

$$D_{\text{erp}}(r) = \frac{\sum_{v_k} C_{v_k}(p_k) D_{v_k}(p_k)}{\sum_{v_k} C_{v_k}(p_k)}. \quad (5)$$

To derive the correlation weight  $C_{v_k}(p_k)$ , we compute three complementary measures: *sharpness*, *locality*, and *symmetry*, from the attention map  $\widetilde{\mathbf{M}}_{\text{Attn}}$ .

- **Sharpness.**  $S_{\text{sharp}}(p_k)$  quantifies the concentration of attention around the point  $p_k$ . A sharply peaked distribution indicates that the model focuses confidently on a few stable correspondences, implying stronger geometric reliability. We compute it using the normalized Shannon entropy, where lower entropy corresponds to sharper attention and thus higher  $S_{\text{sharp}}(p_k)$ :

$$S_{\text{sharp}}(p_k) = 1 - \frac{H(p_k)}{\log|\Omega(v_k)|}, \quad (6)$$

$$H(p_k) = - \sum_{p \in \Omega(v_k)} \widetilde{\mathbf{M}}_{\text{Attn}}(p_k, p) \log \widetilde{\mathbf{M}}_{\text{Attn}}(p_k, p),$$

where  $\Omega(v_k)$  denotes the set of points within view  $v_k$ .

- **Locality.**  $S_{\text{loc}}(p_k)$  measures the spatial compactness of attention around the point  $p_k$ , as stable geometric regions typically attend locally, while long-range attention suggests unreliable local features being compensated by distant cues. We compute  $S_{\text{loc}}(p_k)$  by weighting spatial distances with a Gaussian kernel  $G(\cdot)$ :

$$S_{\text{loc}}(p_k) = \sum_{p \in \Omega(v_k)} \widetilde{\mathbf{M}}_{\text{Attn}}(p_k, p) G(\|\mathbf{x}_p - \mathbf{x}_{p_k}\|), \quad (7)$$

where  $\mathbf{x}_{p_k}$  and  $\mathbf{x}_p$  denote 2D coordinates of  $p_k$  and  $p$ .

Table 1. Quantitative comparisons with SOTA methods on Matterport3D [9], Stanford2D3D [3], and Replica360-2K [27, 30] test sets.

Test Dataset	Method	Backbone	Train $\rightarrow$ Test	Abs Rel $\downarrow$	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$	
Matterport3D	BiFuse++ [37]	ResNet34 [16]	M $\rightarrow$ M	0.112	0.881	0.966	0.987	
	EGFormer [44]	Transformer	M $\rightarrow$ M	0.147	0.816	0.939	0.974	
	HRDFuse [2]	ResNet34 [16]	M $\rightarrow$ M	0.117	0.867	0.962	0.985	
	Elite360D [1]	EfficientNet-B5 [32]	M $\rightarrow$ M	0.105	0.899	0.971	<b>0.991</b>	
	Depth Anywhere [39]	Depth Anything [42], UniFuse [19]	M+ $\rightarrow$ M	0.089	0.911	0.975	<b>0.991</b>	
	Depth Anywhere [39]	Depth Anything [42], BiFuse++ [37]	M+ $\rightarrow$ M	0.085	0.917	0.976	<b>0.991</b>	
	360MD [27]	MiDaS v2 [26]	Training-free	0.264	0.612	0.854	0.941	
	RPG360 [20]	Omnidata v2 [13]	Training-free	0.215	0.796	0.935	0.973	
	RPG360 [20]	Metric3D v2 [17]	Training-free	0.203	0.859	0.953	0.977	
	VGGT-360	VGGT [38]	Training-free	0.083	0.935	0.978	0.989	
	VGGT-360	Fastvggt [29]	Training-free	<b>0.078</b>	<b>0.943</b>	<b>0.981</b>	<b>0.991</b>	
	VGGT-360	$\pi^3$ [40]	Training-free	0.079	0.941	0.979	0.989	
Stanford2D3D	BiFuse [36]	ResNet50 [16]	M $\rightarrow$ S	0.120	0.862	-	-	
	UniFuse [19]	ResNet34 [32]	M $\rightarrow$ S	0.094	0.913	-	-	
	BiFuse++ [37]	ResNet34 [16]	M $\rightarrow$ S	0.107	0.914	0.975	0.989	
	Depth Anywhere [39]	Depth Anything [42], UniFuse [19]	M+ $\rightarrow$ S	0.082	0.927	0.978	0.990	
	Depth Anywhere [39]	Depth Anything [42], BiFuse++ [37]	M+ $\rightarrow$ S	0.083	0.930	0.978	0.990	
	DAC [15]	ResNet101 [16]	In+ $\rightarrow$ S	0.124	0.859	0.976	0.991	
	360MD [27]	MiDaS v2 [26]	Training-free	0.268	0.636	0.878	0.945	
	VGGT-360	VGGT [38]	Training-free	0.068	<b>0.953</b>	0.983	0.991	
	VGGT-360	Fastvggt [29]	Training-free	0.070	0.952	0.983	0.991	
	VGGT-360	$\pi^3$ [40]	Training-free	<b>0.065</b>	0.952	<b>0.984</b>	<b>0.993</b>	
	Replica360-2K	BiFuse [36]	ResNet50 [16]	M $\rightarrow$ R	0.318	0.591	0.840	0.927
		HoHoNet [31]	ResNet50 [16]	M $\rightarrow$ R	0.259	0.672	0.888	0.942
UniFuse [19]		ResNet34 [16]	M $\rightarrow$ R	0.233	0.728	0.905	0.954	
Depth Anywhere [39]		Depth Anything [42], UniFuse [19]	M+ $\rightarrow$ R	0.219	0.763	0.905	0.951	
Depth Anywhere [39]		Depth Anything [42], BiFuse++ [37]	M+ $\rightarrow$ R	0.223	0.758	0.903	0.955	
DAC [15]		ResNet101 [16]	In+ $\rightarrow$ R	0.142	0.803	0.960	0.994	
360MD [27]		MiDaS v2 [26]	Training-free	0.167	0.769	0.948	0.983	
HDE360 [24]		UniFuse [19]	Training-free	0.133	0.870	0.961	0.978	
HDE360 [24]		HoHoNet [31]	Training-free	0.107	0.910	0.961	0.982	
VGGT-360		VGGT [38]	Training-free	0.075	0.934	0.985	0.993	
VGGT-360		Fastvggt [29]	Training-free	<b>0.069</b>	<b>0.950</b>	<b>0.989</b>	<b>0.996</b>	
VGGT-360		$\pi^3$ [40]	Training-free	0.070	0.941	<b>0.989</b>	<b>0.996</b>	

M: Matterport3D [9]    M+: Matterport3D [9] + pseudo-labeled Structured3D [48]    In+: Joint indoor datasets [25, 28, 45]

- **Symmetry.**  $S_{\text{sym}}(p_k)$  quantifies the mutual consistency of attention around the point  $p_k$ , as reliable geometric correspondences are typically bidirectional (*i.e.*, if  $p_k$  attends strongly to another point  $p$ ,  $p$  should also attend back to  $p_k$ ). We compute it using the Bhattacharyya coefficient:

$$S_{\text{sym}}(p_k) = \sum_{u \in \Omega(v_k)} \sqrt{\widetilde{\mathbf{M}}_{\text{Attn}}(p_k, p) \widetilde{\mathbf{M}}'_{\text{Attn}}(p_k, p)}, \quad (8)$$

where  $\widetilde{\mathbf{M}}'_{\text{Attn}}$  is the normalized transpose of  $\widetilde{\mathbf{M}}_{\text{Attn}}$ . The three correlation metrics are additively aggregated and then normalized to derive the correlation weight  $\mathbf{C}_{v_k}(p_k)$ :

$$\mathbf{C}_{v_k}(p_k) = \text{Norm}(\mathbf{S}_{\text{sharp}}(p_k) + \mathbf{S}_{\text{loc}}(p_k) + \mathbf{S}_{\text{sym}}(p_k)), \quad (9)$$

where all metrics are pre-normalized to ensure balanced contributions. As shown in Fig. 6, the refined 3D model with our correction model demonstrates significantly improved geometric coherence and higher accuracy.

## 4. Experiment

**Datasets.** For benchmarking, we evaluate on three standard indoor datasets: **Matterport3D** [9], **Stanford2D3D** [3], and **Replica360-2K** [27, 30], which provide groundtruth depth and cover diverse resolutions and indoor layouts. To assess generalization under outdoor scenarios, following 360MD [27], we further provide qualitative results from **OmniPhotos** [5], where groundtruth is unavailable.

**Implementation.** During the adaptive projection phase, we use  $N_B=8$  base views and augment the top- $K=2$  most uncertain ones. For enhanced attention, we set  $m=0.05$  to modulate the sharpness and border width of the structure-aware confidence map. We evaluate our framework on three VGGT-like baselines: VGGT [38],  $\pi^3$  [40], and Fastvggt [29]. All experiments are conducted on NVIDIA RTX 4090 and TITAN RTX GPUs.

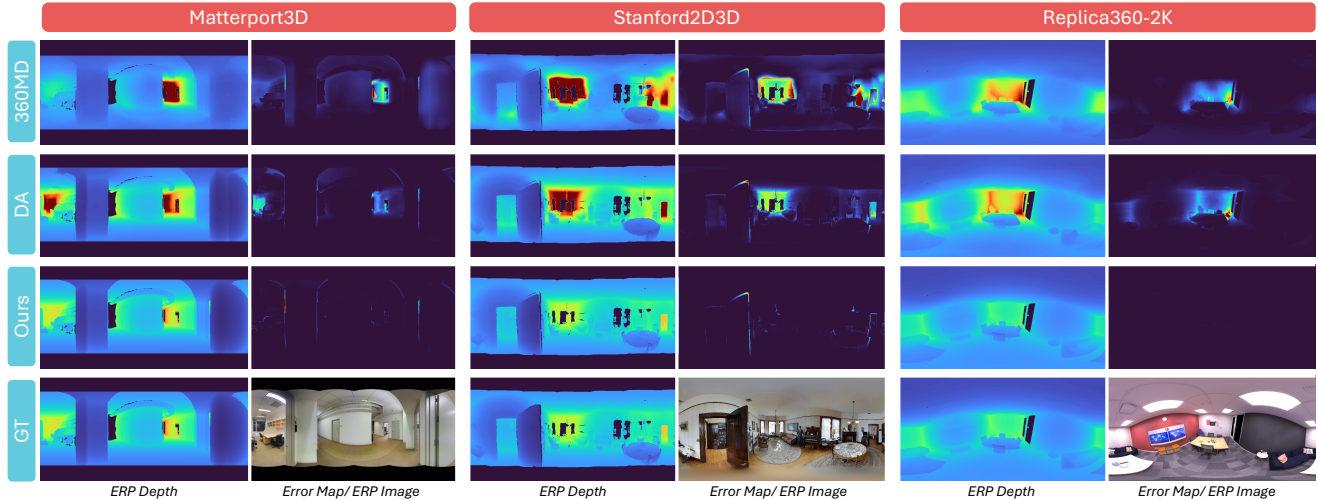


Figure 7. Qualitative comparisons with the state-of-the-art supervised method Depth-Anywhere (DA) [39] and the training-free method 360MD [27] across three indoor datasets: Matterport3D [9], Stanford2D3D [3], and Replica360-2K [27, 30].

**Baselines.** We compare our VGGT-360 with **three** state-of-the-art (SOTA) *training-free* panoramic depth estimation methods: 360MD [27], HDE360 [24], and RPG360 [20], as well as **nine** SOTA *fully-trained* methods: BiFuse [36], UniFuse [19], HoHoNet [31], BiFuse++ [37], EGFormer [44], HRDFuse [2], Elite360D [1], Depth Anywhere [39], and DAC [15]. All fully-trained methods are trained on the Matterport3D training set (**M**), except Depth Anywhere [39], which additionally leverages Structured3D [48] pseudo-labels (**M+**), and DAC [15], which is trained jointly on three indoor datasets [25, 28, 45] (**In+**).

#### 4.1. Comparison with SOTA methods

Table 1 reports quantitative comparisons with SOTA methods across three benchmarks. On **Matterport3D** [9], although all supervised baselines are trained on its official set, VGGT-360 still achieves the best overall performance, outperforming both training-free and fully-trained methods. Notably, it even surpasses Depth Anywhere [39], which is trained on large-scale labeled data and leverages powerful pretrained large depth models. This result highlights the effectiveness of our geometry-aware VGGT-360 framework, which achieves consistent and structure-preserving depth without any task-specific adaptation.

We further assess the zero-shot generalization of all methods on **Stanford2D3D** [3], a challenging benchmark with domain shifts in appearance and layout. While supervised models often rely on dataset-specific priors and texture cues, our VGGT-360 achieves high accuracy by enforcing view-consistent geometric reasoning. This highlights the inherent generalizability of our geometry-grounded design across diverse and unseen environments.

We also conduct zero-shot evaluation on the more challenging **Replica360-2K** [27, 30], which contains high-

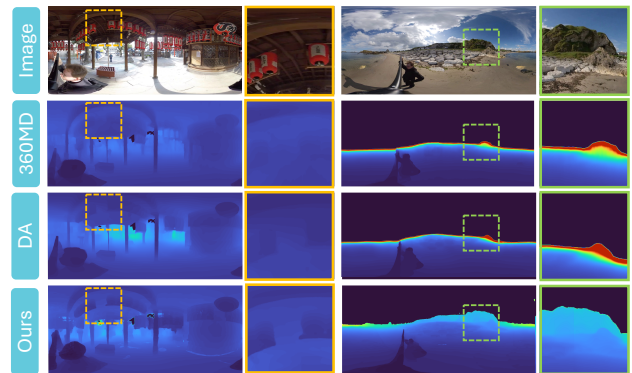


Figure 8. Qualitative comparison with SOTA methods: supervised Depth-Anywhere (DA) [39] and training-free 360MD [27], on outdoor panoramas from OmniPhotos [5].

resolution indoor scenes. Our VGGT-360 still achieves the best overall performance, demonstrating strong scalability to high-resolution inputs. This robustness is attributed to our adaptive projection and multi-view geometric reasoning, which together ensure global consistency and structural fidelity across varying scene complexities and resolutions.

#### 4.2. Visualization on Indoor and Outdoor Scenes

Figure 7 presents qualitative comparisons on **three indoor datasets**: Matterport3D [9], Stanford2D3D [3], and Replica360-2K [27, 30]. We compare our method against a fully-supervised baseline (Depth Anywhere [39]) and a training-free baseline (360MD [27]). Across all scenes, VGGT-360 produces more accurate and globally coherent depth maps, effectively preserving object boundaries and suppressing depth inconsistencies. The advantages of our method are especially evident under zero-shot and high-resolution settings, where baseline approaches often suffer




Figure 9. Ablation studies on Stanford2D3D [3]: effectiveness of the three proposed modules across various VGGT-like baselines.

from over-smoothing and inaccurate geometric predictions. In contrast, our globally consistent pipeline captures scene structure faithfully without training. The accompanying error maps further confirm the superiority of our predictions, showing fewer high-error regions, especially in geometrically challenging areas such as distant surfaces.

Figure 8 presents qualitative comparisons on **outdoor panoramas** from OmniPhotos [5], which features diverse real-world scenes with complex geometry, rich textures, and wide depth ranges. Since no ground-truth depth is available, following 360MD [27], we compare methods based on visual plausibility. Our VGGT-360 yields sharper and more structure-preserving predictions, retaining fine architectural details and geometric consistency. These results highlight the strong zero-shot capability of VGGT-360 in unconstrained outdoor environments, demonstrating its robustness and practicality across diverse panoramic domains.

### 4.3. Ablation Study

**Generalization and Effectiveness of Each Module.** As shown in Fig. 9, all three proposed modules consistently improve accuracy across diverse VGGT-like backbones, validating the strong generality of our design. Each module operates without retraining, exhibits high plug-and-play compatibility across architectures, and collectively delivers robust gains in panoramic depth estimation.

**Impact of Different Projection Strategies.** Building on the above analysis, we conduct ablation on the projection module using Stanford2D3D [3], varying the number of base and adaptively selected neighbor views. As shown in Fig. 10, uniform projection with  $N_B=6$  views offers limited performance, while simply increasing the view count brings marginal gains at a high computational cost. In contrast, our adaptive strategy (*i.e.*,  $N_B=8$  base views with top- $K=2$  neighbor augmentation) achieves a better accuracy–efficiency trade-off, showing that dynamically focusing on uncertain regions outperforms fixed sampling.

**Effect of Structure-Saliency Enhanced Attention.** As shown in Table 2, we analyze the effect of each component in the proposed Structure-Saliency Enhanced At-

Table 2. Ablation of VGGT [38] baseline on Stanford2D3D [3].

Method	Abs Rel↓	RMSE↓	Time
<i>Effect of Structure-Saliency Enhanced Attention</i>			
Baseline	0.080	0.354	<b>1.41s</b>
Baseline + $M_g$	0.075	0.343	1.43s
Baseline + $M_g + E$	0.073	0.346	1.44s
Baseline + $M_g + E + 1_{\text{valid}}$	<b>0.072</b>	<b>0.340</b>	1.45s
<i>Effect of Correlation-Weighted 3D Model Correction</i>			
Baseline	0.080	0.354	<b>1.41s</b>
Baseline + $S_{\text{sharp}}$	0.074	0.328	1.43s
Baseline + $S_{\text{loc}}$	0.073	0.327	1.44s
Baseline + $S_{\text{sym}}$	0.074	0.328	1.44s
Baseline + $S_{\text{sharp}} + S_{\text{loc}} + S_{\text{sym}}$	<b>0.071</b>	<b>0.325</b>	1.46s




Figure 10. Ablation studies on Stanford2D3D [3]: impact of projection parameters (base views ( $N_B=6-12$ ) and top- $K$  neighbor augmentation ( $K=0-3$ )) on performance (a) and runtime (b).

tention. Incorporating the gradient-based saliency prior  $M_g$  improves accuracy by emphasizing structural edges. Adding the edge-band prior  $E$  further enhances boundary regions, while the validity mask  $1_{\text{valid}}$  yields the best results, showing the complementary roles of saliency, edge weighting, and view validity in refining attention.

#### Effect of Correlation-Weighted 3D Model Correction.

We further ablate each component of the correlation-weighted 3D model correction module, as shown in Table 2. Starting from the baseline, adding any single correlation cue (*i.e.*, sharpness  $S_{\text{sharp}}$ , locality  $S_{\text{loc}}$ , or symmetry  $S_{\text{sym}}$ ) improves accuracy. Combining all three yields the best performance, showing that their synergy promotes more reliable 3D structures and higher-quality depth.

## 5. Conclusion

In this paper, we presented VGGT-360, a training-free framework for panoramic depth estimation that reformulated the task as panoramic reprojection from a globally consistent 3D model. Leveraging VGGT’s intrinsic 3D reasoning, our method overcame the limitations of prior view-independent methods without any labels or training. We introduced three key modules: 1) uncertainty-guided adaptive projection for geometry-aware view sampling, 2) structure-saliency enhanced attention for robust 3D reconstruction, and 3) correlation-weighted 3D model correction for reliable ERP depth estimation. Experiments showed VGGT-360 surpassed prior training-free and supervised methods, demonstrating robust zero-shot performance.

## 6. Acknowledgments

This work was supported by the Ministry of Education, Singapore, under its MOE Academic Research Fund Tier 2 (MOE-T2EP20124-0013).

## References

- [1] Hao Ai and Lin Wang. Elite360d: Towards efficient 360 depth estimation via semantic-and distance-aware bi-projection fusion. In *CVPR*, pages 9926–9935, 2024. 6, 7
- [2] Hao Ai, Zidong Cao, Yan-Pei Cao, Ying Shan, and Lin Wang. Hrdfuse: Monocular 360deg depth estimation by collaboratively learning holistic-with-regional depth distributions. In *CVPR*, pages 13273–13282, 2023. 2, 6, 7
- [3] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. 2, 6, 7, 8
- [4] Benjamin Attal, Selenia Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. Matryodshka: Real-time 6dof video view synthesis using multi-sphere images. In *ECCV*, pages 441–459. Springer, 2020. 1
- [5] Tobias Bertel, Mingze Yuan, Reuben Lindroos, and Christian Richardt. Omniphotos: casual 360 vr photography. *ACM Transactions on Graphics (TOG)*, 39(6):1–12, 2020. 6, 7, 8
- [6] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 3
- [7] Zidong Cao and Lin Wang. Crf360d: Monocular 360 depth estimation via spherical fully-connected crfs. *arXiv preprint arXiv:2405.11564*, 2024. 1, 2
- [8] Zidong Cao, Jinjing Zhu, Weiming Zhang, Hao Ai, Haotian Bai, Hengshuang Zhao, and Lin Wang. Panda: Towards panoramic depth anything with unlabeled panoramas and mobius spatial augmentation. In *CVPR*, pages 982–992, 2025. 3
- [9] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *3DV*, 2017. 6, 7
- [10] Yichen Chen, Yuqi Pan, Ruyu Liu, Haoyu Zhang, Guodao Zhang, Bo Sun, and Jianhua Zhang. 360orb-slam: A visual slam system for panoramic images with depth completion network. In *2024 27th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, pages 717–722. IEEE, 2024. 1
- [11] Jiaxi Deng, Yushen Wang, Haitao Meng, Zuoxun Hou, Yi Chang, and Gang Chen. Omnistere: Real-time omnidirectional depth estimation with multiview fisheye cameras. In *CVPR*, pages 1003–1012, 2025. 1
- [12] Marc Eder and Jan-Michael Frahm. Convolutions on spherical images. In *CVPR*, pages 1–5, 2019. 1, 2
- [13] 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 *ICCV*, pages 10786–10796, 2021. 2, 3, 6
- [14] Shijia Ge, Yinxin Zhang, Shuzhao Xie, Weixiang Zhang, Mingcai Zhou, and Zhi Wang. Vggt-dp: Generalizable robot control via vision foundation models. *arXiv preprint arXiv:2509.18778*, 2025. 3
- [15] Yuliang Guo, Sparsh Garg, S Mahdi H Miangoleh, Xinyu Huang, and Liu Ren. Depth any camera: Zero-shot metric depth estimation from any camera. In *CVPR*, pages 26996–27006, 2025. 6, 7
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [17] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 6
- [18] Jingwei Huang, Zhili Chen, Duygu Ceylan, and Hailin Jin. 6-dof vr videos with a single 360-camera. In *2017 IEEE Virtual Reality (VR)*, pages 37–44. IEEE, 2017. 1
- [19] Hualie Jiang, Zhe Sheng, Siyu Zhu, Zilong Dong, and Rui Huang. Unifuse: Unidirectional fusion for 360 panorama depth estimation. *IEEE Robotics and Automation Letters*, 6(2):1519–1526, 2021. 2, 6, 7
- [20] Dongki Jung, Jaehoon Choi, Yonghan Lee, and Dinesh Manocha. Rpg360: Robust 360 depth estimation with perspective foundation models and graph optimization. *arXiv preprint arXiv:2509.23991*, 2025. 2, 3, 6, 7
- [21] Yuyan Li, Yuliang Guo, Zhixin Yan, Xinyu Huang, Ye Duan, and Liu Ren. Omnifusion: 360 monocular depth estimation via geometry-aware fusion. In *CVPR*, pages 2801–2810, 2022. 2
- [22] Yang Liu, Chuanchen Luo, Zimo Tang, Junran Peng, and Zhaoxiang Zhang. Vggt-x: When vggt meets dense novel view synthesis. *arXiv preprint arXiv:2509.25191*, 2025. 3
- [23] Payal Mohadikar, Chuanmao Fan, and Ye Duan. Ms360: A multi-scale feature fusion framework for 360 monocular depth estimation. In *Proceedings of the 50th Graphics Interface Conference*, pages 1–11, 2024. 2
- [24] Chi-Han Peng and Jiayao Zhang. High-resolution depth estimation for 360deg panoramas through perspective and panoramic depth images registration. In *WACV*, pages 3116–3125, 2023. 2, 6, 7
- [25] Santhosh K Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alex Clegg, John Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, et al. Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai. *arXiv preprint arXiv:2109.08238*, 2021. 6, 7
- [26] 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*, 44(3):1623–1637, 2020. 2, 6
- [27] Manuel Rey-Area, Mingze Yuan, and Christian Richardt. 360monodepth: High-resolution 360deg monocular depth estimation. In *CVPR*, pages 3762–3772, 2022. 1, 2, 3, 6, 7, 8

- [28] 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 *ICCV*, pages 10912–10922, 2021. 6, 7
- [29] You Shen, Zhipeng Zhang, Yansong Qu, and Liujuan Cao. Fastvggt: Training-free acceleration of visual geometry transformer. *arXiv preprint arXiv:2509.02560*, 2025. 2, 3, 6
- [30] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 2, 6, 7
- [31] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Hohonet: 360 indoor holistic understanding with latent horizontal features. In *CVPR*, pages 2573–2582, 2021. 2, 6, 7
- [32] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114, 2019. 6
- [33] Keisuke Tateno, Nassir Navab, and Federico Tombari. Distortion-aware convolutional filters for dense prediction in panoramic images. In *ECCV*, pages 707–722, 2018. 2
- [34] An Dinh Vuong, Minh Nhat Vu, and Ian Reid. Improving robotic manipulation with efficient geometry-aware vision encoder. *arXiv preprint arXiv:2509.15880*, 2025. 3
- [35] Fu-En Wang, Hou-Ning Hu, Hsien-Tzu Cheng, Juan-Ting Lin, Shang-Ta Yang, Meng-Li Shih, Hung-Kuo Chu, and Min Sun. Self-supervised learning of depth and camera motion from 360 videos. In *ACCV*, pages 53–68. Springer, 2018. 2
- [36] Fu-En Wang, Yu-Hsuan Yeh, Min Sun, Wei-Chen Chiu, and Yi-Hsuan Tsai. Bifuse: Monocular 360 depth estimation via bi-projection fusion. In *CVPR*, pages 462–471, 2020. 3, 6, 7
- [37] Fu-En Wang, Yu-Hsuan Yeh, Yi-Hsuan Tsai, Wei-Chen Chiu, and Min Sun. Bifuse++: Self-supervised and efficient bi-projection fusion for 360 depth estimation. *IEEE transactions on pattern analysis and machine intelligence*, 45(5): 5448–5460, 2022. 2, 6, 7
- [38] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *CVPR*, pages 5294–5306, 2025. 2, 3, 6, 8
- [39] Ning-Hsu Albert Wang and Yu-Lun Liu. Depth anywhere: Enhancing 360 monocular depth estimation via perspective distillation and unlabeled data augmentation. *NeurIPS*, 37: 127739–127764, 2024. 1, 2, 3, 6, 7
- [40] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He.  $\pi^3$ : Scalable permutation-equivariant visual geometry learning. *arXiv preprint arXiv:2507.13347*, 2025. 2, 3, 6
- [41] Zhiqiang Yan, Xiang Li, Kun Wang, Shuo Chen, Jun Li, and Jian Yang. Distortion and uncertainty aware loss for panoramic depth completion. In *ICML*, pages 39099–39109. PMLR, 2023. 1
- [42] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, pages 10371–10381, 2024. 2, 3, 6
- [43] 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 *ICCV*, pages 9043–9053, 2023. 2
- [44] Ilwi Yun, Chanyong Shin, Hyunku Lee, Hyuk-Jae Lee, and Chae Eun Rhee. Egformer: Equirectangular geometry-biased transformer for 360 depth estimation. In *ICCV*, pages 6101–6112, 2023. 1, 2, 6, 7
- [45] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, pages 3712–3722, 2018. 6, 7
- [46] Junsong Zhang, Zisong Chen, Chunyu Lin, Zhijie Shen, Lang Nie, Kang Liao, and Yao Zhao. Sgformer: Spherical geometry transformer for 360 depth estimation. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025. 1, 2
- [47] Yi Zhang and Fei Huang. Panoramic visual slam technology for spherical images. *Sensors*, 21(3):705, 2021. 1
- [48] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. In *ECCV*, pages 519–535. Springer, 2020. 6, 7
- [49] Chuanqing Zhuang, Zhengda Lu, Yiqun Wang, Jun Xiao, and Ying Wang. Acdnet: Adaptively combined dilated convolution for monocular panorama depth estimation. In *AAAI*, pages 3653–3661, 2022. 2