

360MonoDepth: High-Resolution 360° Monocular Depth Estimation

Manuel Rey-Area* Mingze Yuan* Christian Richardt
University of Bath

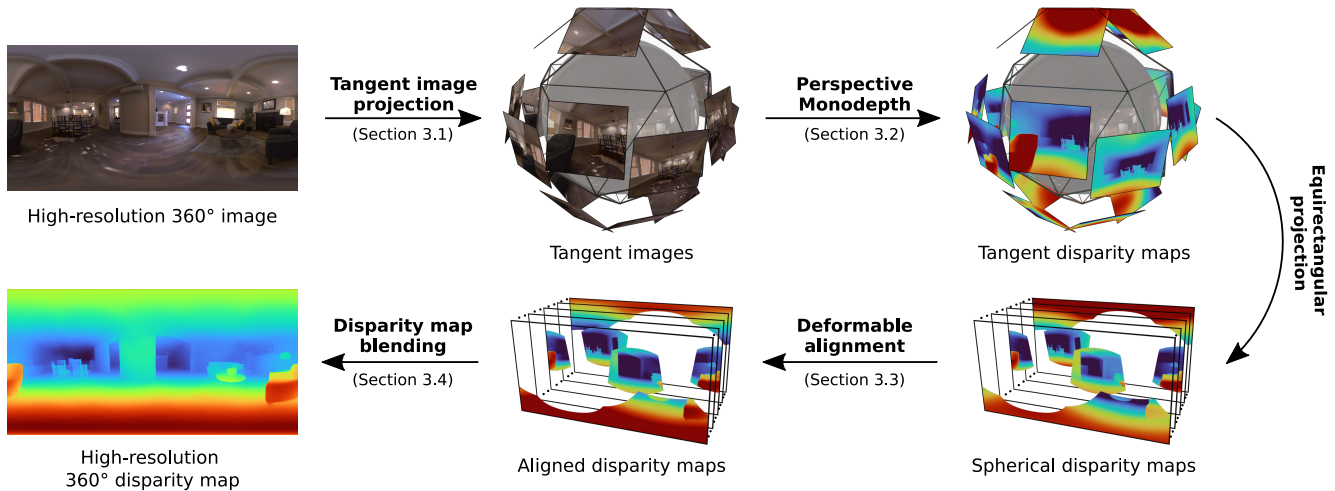


Figure 1. We present a flexible framework for estimating high-resolution disparity maps from a single 360° input image by decomposing it into perspective tangent images, which are used for monocular depth estimation. We then globally align all disparity maps using multi-scale alignment fields, and blend them in the gradient domain to produce a detailed, consistent and high-resolution 360° spherical disparity map.

Abstract

360° cameras can capture complete environments in a single shot, which makes 360° imagery alluring in many computer vision tasks. However, monocular depth estimation remains a challenge for 360° data, particularly for high resolutions like 2K (2048×1024) and beyond that are important for novel-view synthesis and virtual reality applications. Current CNN-based methods do not support such high resolutions due to limited GPU memory. In this work, we propose a flexible framework for monocular depth estimation from high-resolution 360° images using tangent images. We project the 360° input image onto a set of tangent planes that produce perspective views, which are suitable for the latest, most accurate state-of-the-art perspective monocular depth estimators. To achieve globally consistent disparity estimates, we recombine the individual depth estimates using deformable multi-scale alignment followed by gradient-domain blending. The result is a dense, high-resolution 360° depth map with a high level of detail, also for outdoor scenes which are not supported by existing methods. Our source code and data are available at <https://manurare.github.io/360monodepth/>.

1. Introduction

Monocular depth estimation has recently seen a significant boost thanks to convolutional neural networks. CNNs have demonstrated an unprecedented expressive power to learn intricate geometric relationships from data, resembling the capability of humans to exploit visual cues to perceive depth. Monocular depth estimates have enabled impressive new approaches for 3D photography [33, 61] and novel-view synthesis of dynamic scenes [20, 43]. However, most approaches for monocular depth estimation are limited to low-resolution¹ perspective images, with a limited field-of-view.

Nevertheless, 360° cameras are becoming increasingly popular and widespread in the computer vision community. The omnidirectional 360° field-of-view captured by these devices is appealing for tasks such as robust, omnidirectional SLAM [66, 77], scene understanding and layout estimation [31, 67, 75, 81], or VR photography and video [5, 59]. State-

*These authors contributed equally to this work.

¹For example, 384×384≈0.15 megapixels for MiDaS [55, 56].

of-the-art monocular depth estimation approaches for 360° images [30, 40, 52, 67, 74] are currently limited to resolutions of $1024 \times 512 \approx 0.5$ megapixels. While this is sufficient for tasks like layout estimation, it is insufficient for VR applications as they require resolutions of at least 2 megapixels to match the resolution of VR headsets [34] and achieve full immersion [12, 45]. Our work aims to fill this gap.

Existing monocular 360° depth estimation approaches build on CNNs whose spatial resolution is fundamentally limited by the GPU memory available during training. These methods are therefore restricted to small batch sizes of 4 to 8 for 0.5 megapixel images on an NVIDIA 2080 Ti with 11 GB memory [30, 52, 67]. For this reason, single-CNN approaches become impractical for predicting high-resolution depth maps with multiple megapixels.

In this work, we introduce a general and flexible framework for monocular depth estimation from high-resolution 360° images inspired by Eder et al.’s tangent images [16]. Our approach projects the input 360° image to a collection of perspective tangent images, e.g. using the faces of an icosahedron. We then use state-of-the-art perspective monocular depth estimators endowed with powerful generalisation capability for obtaining dense, detailed depth maps for each tangent image. Subsequently, we optimally align individual depth maps using multi-scale spatially-varying deformation fields to bring them into global agreement. Finally, we merge the aligned depth maps using gradient-based blending for a seamless high-resolution 360° depth map. Our technical contributions are as follow:

1. A simple, yet powerful and practical framework for high-quality multi-megapixel 360° monocular depth estimation based on aligning and blending depth maps predicted from perspective tangent images.
2. Support for increased resolutions using tangent images, and improved quality by forward compatibility for future monocular depth estimation approaches.
3. We provide 2048×1024 ground-truth depth maps for Matterport3D’s stitched skyboxes to advance future high-resolution depth estimation approaches.

2. Related Work

Monocular depth estimation. Predicting a dense depth map from a single input image is a challenging, ill-posed task due to the high level of ambiguity between possible reconstructions. Early approaches relied on simple geometric assumptions [27], geometric reasoning using Markov random fields [58], or non-parametric depth transfer [32]. The rise of deep learning has made it possible to train convolutional neural networks that are supervised by ground-truth depth maps [17, 36, 44], e.g. from synthetic renderings or depth sensors, or by exploiting defocus blur [60, 62]. However, suitable training data is scarce, particularly for outdoor scenes.

Subsequent work therefore explored alternative training regimes, in particular from stereo views that provide self-supervision via view synthesis [21, 22, 23, 47, 54, 72, 78], from camera ego-motion in videos [24, 46, 48, 57, 71, 82, 85], and from multiview stereo reconstructions [41, 42]. Ranftl and Lasinger et al.’s MiDaS [56] demonstrated substantial improvements and generalisation performance by learning from five varied datasets using multi-objective learning. The fidelity of depth predictions can also be improved by merging estimates at multiple scales [49]. Recently, Ranftl et al. [55] introduced transformers [14, 70] into monocular depth estimation, to produce finer-grained and more globally consistent results than CNN-based methods. We base our new monocular 360° depth estimation method on their state-of-the-art performance, but our method would transparently benefit from future advances in monocular depth estimation.

Spherical CNNs. Most CNNs are applied to flat 2D images with little image distortion. However, 360° images need a different approach to correctly handle the inevitable distortions of projecting a spherical image onto a plane, e.g. in the commonly used equirectangular projection. Su and Grauman proposed a pragmatic solution using wider kernels near the poles [64]. However, these kernels do not share any information, which leads to suboptimal performance. Another pragmatic approach is to project the spherical image into a padded cubemap, process all sides as perspective images, and to recombine the results [8]. This approach struggles for the top and bottom faces, as kernel orientations become ambiguous due to 90-degree rotational symmetry. Eder et al. [16] generalise this approach to more than six tangent images, which achieves higher and more uniform angular pixel resolutions. However, predictions on tangent images are recombined per pixel without any alignment or blending, which works poorly for monocular depth estimation (see our experiments in Section 4.3).

Cubemaps have since been generalised to the 20 triangular faces of an icosahedron, which can be unwrapped into 5 rectangles with shared convolution kernels [10, 37, 83]. Distortion-aware convolutions [11, 19, 65, 69, 84] can directly model the distortions of equirectangular projection. Interestingly, this also enables the transfer of models trained on perspective images to equirectangular images without any additional training, but it requires matching angular pixel resolutions. Full rotation-equivariance can be achieved using spherical convolutions [9, 18], but this may not always be desirable as the down direction is usually consistent with gravity. These approaches have high memory requirements that make them unsuitable for multi-megapixel resolutions.

360° depth estimation. Deep learning has also boosted monocular depth estimation for 360° images. Most methods are supervised using synthetic datasets due to the difficulty of acquiring ground-truth spherical depth maps [15, 87]. Similar to the perspective case, several methods perform

self-supervised training via view synthesis [39, 51, 73, 88]. Tateno et al. [69] adapt pre-trained monocular depth estimation for perspective images [36] to spherical images using distortion-aware convolutional filters. Depth accuracy can be improved by fusing predictions for equirectangular and cubemap projections [4, 74], while deformable [7] or dilated [86] convolutions can make methods more distortion-aware. Pintore et al. [52] and Sun et al. [67] exploit gravity-aligned features in man-made interior environments using vertical slicing. However, the performance of these learning-based approaches highly depends on their training data. Most datasets are synthetic, low-resolution (1024×512) and only consider indoor scenes. These methods therefore tend to perform poorly on real high-resolution or outdoor scenes.

Learning-based spherical stereo methods again mostly rely on synthetic training data, making them unsuitable for real outdoor scenes. They assume a known, fixed camera baseline [35, 38, 76], or estimate the relative pose between cameras [73]. Under the assumption of a moving camera in a static environment, structure-from-motion and multi-view stereo can be used [13, 28, 29]. However, these assumptions are violated by most usage scenarios, in which the camera might be stationary or environments are dynamic. Crucially, these techniques do not work for a single monocular input image as information from multiple viewpoints or points in time must be combined.

3. The 360MonoDepth framework

Our approach builds on a general framework for estimating high-resolution depth maps from just a single monocular 360° input image. Figure 1 illustrates the four main steps of our approach. We start by projecting the 360° input image to a set of overlapping perspective tangent images (Section 3.1), for instance the 20 faces of an icosahedron for an equirectangular image of resolution 2048×1024 pixels. For each tangent image, we independently predict a depth map (Section 3.2) using state-of-the-art perspective monocular depth estimation [55, 56]. Such methods predict disparity maps that are ambiguous up to affine ambiguity with unknown scale and shift [79]. We thus formulate a global optimisation to align all tangent disparity maps in the spherical domain (Section 3.3). Finally, we merge the aligned tangent disparity maps using Poisson blending [53] into a high-resolution spherical disparity map (Section 3.4).

In this paper, we use equirectangular projection (ERP) as the default format for spherical 360° images due to its wide adoption in the computer vision community. However, our approach can easily be adapted to any other spherical projection by adapting the projection to/from tangent images.

3.1. Tangent image projection

Carl Friedrich Gauss proved that any projection of a spherical image to a plane introduces some degree of distortion.




Figure 2. Coverage of the sphere by the 20 tangent images of an icosahedron (with padding factor $p = 0.3$). The darkest regions have an overlap of 2, the brightest of 5 images.

For example, equirectangular projection stretches the regions near the poles across the longitudinal dimension. To minimise distortion, we project the spherical image to a set of perspective tangent images, each of which can be processed separately and then recombined. We found it convenient to work with the 20 tangent images produced by the faces of an icosahedron that circumscribes a sphere, as this arrangement fairly uniformly covers the sphere’s surface (see Figure 2), but our framework easily adapts to different numbers. Each triangular face of the icosahedron is tangent to the sphere at its centroid, which we use to create the tangent images using gnomonic projection.

Padding. By default, the size of each tangent image is constrained by the size of its icosahedron face, producing a field of view of 72° . Tightly cropped tangent images include some overlap with adjacent icosahedron faces that share an edge, by nature of packing a triangular shape into a rectangular image (see the blue region in Figure 3). However, more overlap between tangent images, especially for icosahedron faces that only share a single vertex, is desirable for providing consistency constraints in our disparity map alignment step in Section 3.3, as this helps find a globally consistent alignment. Therefore, we extend the boundaries of tangent images by a padding factor of $p \in [0, 1]$ relative to the base shape, as illustrated in Figure 3. We use a padding of $p = 0.3$, which extends the default tangent image by 30% in all directions.




Figure 3. Each icosahedron face (thick triangle outline) is fit within a rectangular tangent image (blue) without padding, i.e. $p = 0$. The green region shows a padding of $p = 0.1$, and red shows $p = 0.2$. Right: Equirectangular projection for two padded tangent images.

3.2. Tangent disparity map estimation

We use monocular depth estimation on each individual tangent image to predict dense disparity maps that will be aligned and merged in the next steps. Specifically, we use MiDaS v2 [56] and v3 [55] for their state-of-the-art performance for both indoor and outdoor images. Nevertheless, our framework is agnostic to the specific perspective monocular depth estimator and will benefit from future improvements.

MiDaS predicts disparity maps that correspond to inverse depth, but with an unknown scale factor and shift offset due to its scale- and shift-invariant training procedure. Our method works consistently in disparity space, as this improves the numerical stability during the optimisation in Section 3.3, particularly for distant parts of the environment.

Perspective to spherical disparity. Perspective disparity maps, as predicted by MiDaS, describe disparity estimates with respect to the viewing direction of a tangent image, i.e. the z -component of a camera ray to a 3D point (in camera coordinates). However, each tangent image has a different viewing direction, so the definitions of disparity are incompatible between tangent images. In contrast, spherical disparity is the inverse (radial) Euclidean distance from the camera’s centre of projection to a 3D point. This definition is consistent for all tangent images as they all share the same centre of projection. We convert the tangent disparity maps from perspective to spherical disparity, and from tangent image space to the equirectangular projection of the input image in preparation for the disparity map alignment step.

3.3. Global disparity map alignment

The individual disparity maps $D(\cdot)$ estimated in the previous step may have inconsistent scales and offsets, as they are predicted independently from each other. Nonetheless, each individual prediction should by design correspond to the ground-truth disparity (i.e. inverse depth) subject to a different unknown affine transform (i.e. scale and offset). To ensure that disparity estimates are consistent with each other, we need to align them globally by finding suitable scale and offset values for each disparity map.

Our global disparity map alignment method is inspired by Hedman and Kopf’s deformable depth alignment [26]. Instead of finding a constant scale and offset per disparity map, they use spatially varying affine adjustment fields. These adjustment fields are modelled as 2D grids of size $m \times n$ in tangent image space. Each grid-point i stores a pair of scale and offset variables (s^i, o^i) that are interpolated bilinearly across the tangent image domain. The rescaled disparity \tilde{D} of a pixel at position \mathbf{x} is computed using

$$\tilde{D}(\mathbf{x}) = s(\mathbf{x})D(\mathbf{x}) + o(\mathbf{x}), \quad (1)$$

where $s(\mathbf{x}) = \sum_i w_i(\mathbf{x})s^i$ and $o(\mathbf{x}) = \sum_i w_i(\mathbf{x})o^i$ are the interpolated scale and offset values, and $w_i(\mathbf{x})$ the bilinear interpolation weights for pixel location \mathbf{x} .

To globally align all tangent disparity maps, we optimise for the affine adjustment fields that minimise the energy

$$\operatorname{argmin}_{\{s_a^i, o_a^i\}} E_{\text{alignment}} + \lambda_{\text{smoothness}} E_{\text{smoothness}} + \lambda_{\text{scale}} E_{\text{scale}}, \quad (2)$$

which trades off alignment with the spatial smoothness of adjustment fields and a scale regularisation term. We use $\lambda_{\text{smoothness}} = 40$ and $\lambda_{\text{scale}} = 0.007$ for all results.

Disparity alignment term. Once aligned, disparity maps should agree where they overlap as they represent the same region of a scene. Given the set \mathcal{T} of tangent image indices, we create the set $\mathcal{Z} = \{(a, b) \mid a, b \in \mathcal{T}, a < b\}$ of ordered pairs of tangent images and use $\Omega(a, b)$ to denote the set of overlapping pixels in images a and b . We quantify the alignment between rescaled disparity maps \tilde{D}_a and \tilde{D}_b using:

$$E_{\text{alignment}} = \frac{1}{z_a} \sum_{(a,b) \in \mathcal{Z}} \sum_{\mathbf{x} \in \Omega(a,b)} (\tilde{D}_a(\mathbf{x}) - \tilde{D}_b(\mathbf{x}))^2, \quad (3)$$

where $z_a = \sum_{(a,b) \in \mathcal{Z}} |\Omega(a, b)|$ is used for normalising by the number of considered pixel pairs. For efficiency, we only sample 1% of pixels from the overlap regions $\Omega(a, b)$.

Smoothness term. We encourage the deformable adjustment fields to be spatially smooth between neighbouring grid-points i and j using

$$E_{\text{smoothness}} = \frac{1}{z_s} \sum_{a \in \mathcal{T}} \sum_{(i,j)} \|s_a^i - s_a^j\|_2^2 + \|o_a^i - o_a^j\|_2^2, \quad (4)$$

where $z_s = |\mathcal{T}| \cdot m \cdot n$ normalises by the number of grid-points in all tangent images.

Scale term. The final term regularises the scale to avoid a collapse to the trivial solution of scale $s = 0$:

$$E_{\text{scale}} = \sum_{a \in \mathcal{T}} \sum_i (s_a^i)^{-1}. \quad (5)$$

Initialisation. We standardise the input spherical disparity maps to unit scale and zero offset [56] using

$$D'(\mathbf{x}) = \frac{D(\mathbf{x}) - \operatorname{median}(D)}{|\mathcal{P}|^{-1} \sum_{\mathbf{x} \in \mathcal{P}} |D(\mathbf{x}) - \operatorname{median}(D)|} \quad (6)$$

to pre-align their ranges, where \mathcal{P} is the set of pixel coordinates. Similarly, we initialise the deformation fields to unit scale $s_a^i = 1$ and zero offset $o_a^i = 0$ for all a and i .

3.3.1 Multi-scale deformable alignment

Different from Hedman and Kopf, we perform deformable alignment at multiple scales, which we found to be beneficial for fine-tuning the global alignment. We start by optimising for a coarse deformation grid of 4×3 grid-points per tangent disparity map. We then apply these deformation fields to the disparity maps, and perform a new optimisation for a 8×7 grid without re-standardising the input disparity maps. We again apply these deformation fields to the disparity maps, and perform a final refinement with a grid size of 16×14 .

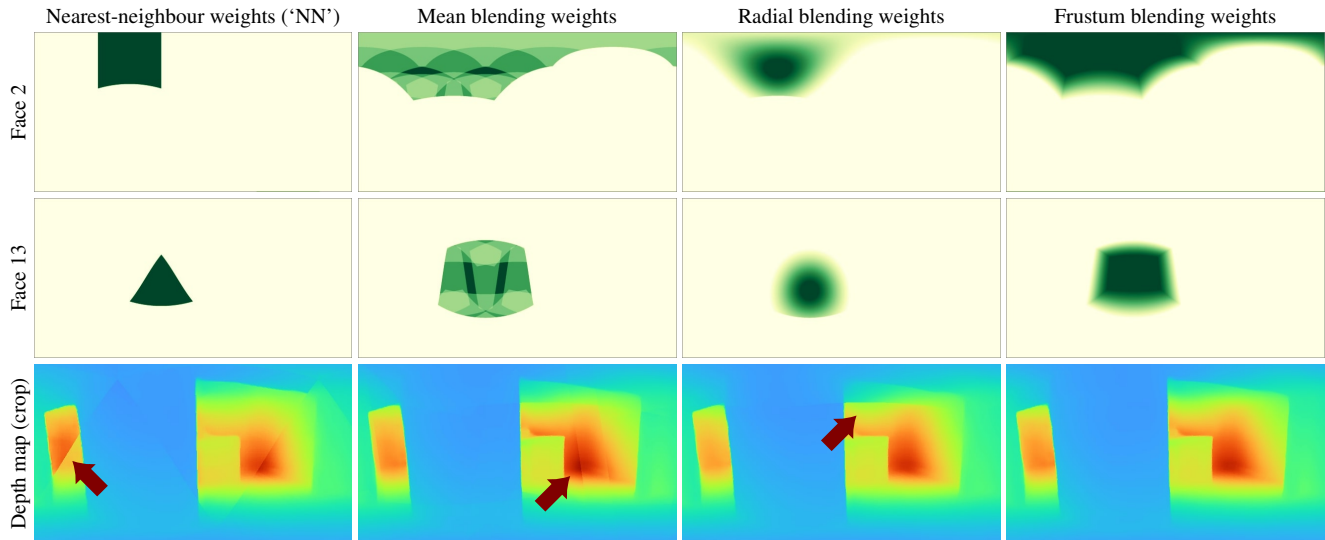


Figure 4. Comparison of blending weights for icosahedron tangent images, in equirectangular projection. Vanilla tangent images [16] select estimates only from the nearest tangent image ('NN'). Mean weights average all overlapping tangent images per pixel. Radial weights start decaying at 15° from the centre of projection. Frustum blending weights start decaying 30% diagonally towards the principal point off from each corner. Notice that disparity maps blended using 'NN', 'mean' and 'radial' weights contain visible seams, which 'frustum' minimises.

3.4. Disparity map blending

After the alignment, the individual disparity maps need to be merged into a single spherical disparity map, similar to how multiple photos are merged into a panorama during stitching. Naïvely merging the tangent disparity maps using nearest-neighbour ('NN') or averaging per-pixel ('mean') leads to undesirable seams, as shown in Figure 4. Using smoothly feathered blending weights [80] in the shape of a frustum reduces seams, but may produce blurrier results.

For the highest fidelity blending, we take inspiration from panorama stitching [68] and blend disparity maps in the gradient domain using Poisson blending [53]. Specifically, we look for the blended disparity map $B(\cdot)$ that minimises:

$$\begin{aligned} \operatorname{argmin}_B \sum_{a \in \mathcal{T}} \sum_{\mathbf{x}} \omega_a(\mathbf{x}) \left\| \nabla B(\mathbf{x}) - \nabla \tilde{D}_a(\mathbf{x}) \right\|_2^2 \\ + \lambda_{\text{fidelity}} \cdot \sum_{\mathbf{x}} (B(\mathbf{x}) - D_{\text{NN}}(\mathbf{x}))^2 \end{aligned} \quad (7)$$

where $\omega_a(\mathbf{x})$ are the spatially varying 'frustum' blending weights that modulate the influence of pixels (see Figure 4), and $\lambda_{\text{fidelity}} = 0.1$ is a weight to encourage the solution to stay close to the nearest-neighbour disparity map stitch D_{NN} .

4. Experiments and Results

Implementation. When processing equirectangular images at a resolution of 2048×1024 pixels, we use the 20 tangent images of a icosahedron. We project each tangent image using a padding of $p = 0.3$ to a resolution of 400×346 pixels. This closely matches the 384×384 training resolution used by MiDaS v2/v3 [55, 56], for which we use the authors'

implementation. We solve the global disparity map alignment problem in Equation 2 using the Ceres non-linear least-squares solver [1]. Specifically, we perform L-BFGS line search for 50 iterations at each scale. The gradient-based disparity map blending in Equation 7 is a large sparse least-squares problem that we solve using Eigen's biconjugate gradient stabilized solver (BiCGSTAB) [25]. As the Matterport3D dataset [6] does not include the top and bottom regions of the scene, we exclude a circular region of radius 25° at the top and bottom from our alignment step.

Datasets. For benchmarking, we use equirectangular input images and ground-truth depth maps created from the Matterport3D [6] and Replica [63] datasets. These datasets contain indoor environments reconstructed as a textured mesh and thus provide ground-truth depth. We also show qualitative results on varied outdoor images from OmniPhotos [5], for which no ground-truth depth maps are available.

Matterport3D [6] is a real indoor dataset that comprises 10,800 panoramic images. Unfortunately, the poses of these 'skybox' images relative to the mesh reconstruction are not provided, which prevents rendering aligned ground-truth depth maps. Previous work overcame this by rendering both images and depth maps from the textured mesh [87]. However, the image quality of these synthetic images is worse than the real skybox images, particularly at the 2048×1024 resolution we are targeting. We therefore estimate the poses for the real skybox images relative to the mesh using 360° structure-from-motion [50] applied to a mixture of real and rendered skybox images at known camera positions. The estimated camera poses allow us to render ground-truth depth maps with pixel accuracy from the provided scene mesh.

Table 1. Quantitative results for Matterport3D-2K and Replica360-2K, at 2048×1024 with Poisson blending. Highlighting: **best**, **second-best**.

Method	Matterport3D-2K							Replica360-2K						
	AbsRel▼	MAE▼	RMSE▼	RMSE-log▼	$\delta < 1.25^\Delta$	$\delta < 1.25^2^\Delta$	$\delta < 1.25^3^\Delta$	AbsRel▼	MAE▼	RMSE▼	RMSE-log▼	$\delta < 1.25^\Delta$	$\delta < 1.25^2^\Delta$	$\delta < 1.25^3^\Delta$
OmniDepth [87]	0.473	0.946	1.317	0.212	0.378	0.647	0.820	0.352	0.589	0.787	0.168	0.479	0.776	0.906
BiFuse [74]	0.321	0.649	0.994	0.158	0.564	0.802	0.910	0.318	0.468	0.663	0.152	0.591	0.840	0.927
HoHoNet ^M [67]	0.227	0.430	0.686	0.132	0.723	0.887	0.946	0.259	0.381	0.520	0.131	0.672	0.888	0.942
HoHoNet ^S [67]	0.234	0.487	0.736	0.120	0.654	0.886	0.959	0.221	0.355	0.480	0.112	0.701	0.905	0.960
UniFuse ^M [30]	0.200	0.396	0.652	0.113	0.769	0.908	0.958	0.233	0.330	0.474	0.120	0.728	0.905	0.954
Ours ^{M2} (single-scale)	0.223	0.491	0.828	0.129	0.619	0.867	0.953	0.182	0.412	0.732	0.095	0.750	0.935	0.971
Ours ^{M3} (single-scale)	0.210	0.476	0.840	0.121	0.656	0.889	0.958	0.192	0.447	0.805	0.100	0.737	0.925	0.969
Ours ^{M2} (multi-scale)	0.224	0.494	0.831	0.130	0.616	0.866	0.953	0.167	0.364	0.619	0.089	0.769	0.948	0.981
Ours ^{M3} (multi-scale)	0.208	0.446	0.791	0.119	0.656	0.890	0.961	0.198	0.465	0.841	0.103	0.730	0.920	0.965

^M Trained on Matterport3D [6] ^S Trained on Stanford 2D-3D-S [2] ^{M2} Using MiDaS v2 [56] ^{M3} Using MiDaS v3 [55]

From the original test split of Matterport3D with 2,014 samples, we managed to estimate accurate camera poses for 1,850 (92%) skybox images, and rendered the aligned ground-truth depth maps at 2048×1024 resolution. We will make skybox poses and ground-truth depth maps available.

To assess the generalisation capability and scalability of our framework against baselines, we also evaluate on 360° RGBD data from the Replica dataset [63], which features high-quality indoor room scans that have not been used for training any method. For 13 rooms, we rendered 10 images and ground-truth depth maps at 2048×1024 and 4096×2048 resolution with random poses using the Replica360 renderer [3], for a total of 130 samples each.

Baselines. We compare our results to OmniDepth [87], BiFuse [74], HoHoNet [67] and UniFuse [30] using the authors’ public implementations and pretrained weights. OmniDepth is trained for 512×256 input, while the other methods are for 1024×512. For each method, we downscale the input images to match the expected resolution, and upsample the estimated depth map bilinearly to the input image resolution.

Metrics. We use the standard evaluation metrics adopted for monocular depth estimation evaluation [17]. Although our method operates in disparity space, we report metrics in depth space for fair comparisons with baselines. Please see our supplemental document for details.

4.1. Quantitative evaluation

Table 1 shows the quantitative comparison of our method to the baselines on the Matterport3D-2K and Replica360-2K test sets. Matterport3D is often used for training and evaluating 360° monodepth methods. Indeed, methods trained on it (HoHoNet, UniFuse) tend to perform best. Our method produces competitive results (in several metrics) without any training on Matterport3D, while producing depth maps at a higher resolution and level of detail (see Figures 5 and 6). Replica360 has not been used for training any method, so we can use it to measure generalisation to unseen data. In

Table 2. Quantitative results for Replica360-4K at 4096×2048 with frustum blending (best trade-off between runtime and performance). For superscripts, see Table 1. Highlighting: **best**, **second-best**.

Method	AbsRel▼	MAE▼	RMSE▼	RMSE-log▼	$\delta < 1.25^\Delta$	$\delta < 1.25^2^\Delta$	$\delta < 1.25^3^\Delta$
OmniDepth	0.337	0.582	0.778	0.161	0.484	0.785	0.920
BiFuse	0.292	0.445	0.637	0.143	0.606	0.857	0.941
HoHoNet ^M	0.251	0.379	0.509	0.127	0.670	0.884	0.948
HoHoNet ^S	0.208	0.335	0.455	0.106	0.728	0.909	0.961
UniFuse ^M	0.223	0.324	0.464	0.116	0.744	0.910	0.959
Ours ^{M2} (multi-scale)	0.150	0.335	0.558	0.081	0.813	0.953	0.983
Ours ^{M3} (multi-scale)	0.161	0.363	0.607	0.085	0.781	0.951	0.984

most metrics, our approach clearly outperforms the baselines, which struggle to generalise to this new dataset. The other two metrics, MAE and RMSE, are closely related to the L1 and BerHu (mixed L1/L2) losses used for training HoHoNet [67] and UniFuse [30], respectively, which explains these methods’ better performance in these specific metrics. We further show results at 4K resolution in Table 2. Our results improved across all metrics compared to 2K resolution, and our approach ranks as top-2 in 6 out of 7 metrics, up from 5 out of 7 at 2K resolution (8% improvement in MAE). This shows that our method robustly scales to higher resolutions.

4.2. Qualitative comparisons

We show qualitative comparisons in Figure 5, 6 and 7, and our supplemental results website. For datasets with available ground-truth depth maps, we show depth maps, otherwise disparity maps. On Matterport3D, our results are mostly on par with UniFuse (best in Table 1). On Replica360, our results show fewer errors and cleaner surfaces. Our approach clearly outperforms the baselines on the outdoor OmniPhotos, as no baseline is trained on outdoor data. Our results show the highest level of detail and the sharpest depth edges.

4.3. Ablation studies

We perform two ablation studies to test our design choices in the disparity maps alignment and blending stages of our




Figure 5. Qualitative comparison to different methods on different datasets. Our results show the highest level of detail of all predictions.

method, summarised in Table 3. Our multi-scale alignment and Poisson blending approaches outperform other alternatives. In particular, our alignment step substantially outperforms the “No alignment” of Eder et al. [16] across all metrics. Both deformable multi-scale alignment and blending are necessary for the best results.

5. Discussion and Conclusion

Our method can fail if the tangent disparity estimates are incorrect, e.g. for large plain walls, saturated skies, or photo-realistic wallpapers. As these estimates improve over time, our method can take advantage of them. In some cases, the least-squares rescaling to fit the ground-truth disparity results in negative disparities, which produces incorrect, negative depth values. We also saw inconsistencies in the ground-truth depth maps, such as mirrors or missing lamps or chandeliers that are visible in the image. We show examples of these failure cases in the supplemental document.

Table 3. Ablation studies for disparity map alignment (top) and blending (bottom), evaluated on the Matterport3D test set. Multi-scale deformable alignment outperforms all single-scale alignments across all metrics when using MiDaS v3. Gradient-based Poisson blending outperforms simpler blending modes in all but one metric when using MiDaS v2. Highlighting: **best**, **second-best**.

Method	AbsRel ∇	MAE ∇	RMSE ∇	RMSE-log ∇	$\delta < 1.25^{\blacktriangle}$	$\delta < 1.25^2^{\blacktriangle}$	$\delta < 1.25^3^{\blacktriangle}$
No alignment ^{M3}	0.259	0.600	0.969	0.150	0.532	0.821	0.933
2×2 single-scale ^{M3}	0.210	0.476	0.838	0.122	0.654	0.888	0.959
4×3 single-scale ^{M3}	0.210	0.475	0.838	0.121	0.655	0.889	0.959
8×7 single-scale ^{M3}	0.210	0.476	0.840	0.121	0.656	0.889	0.958
16×14 single-scale ^{M3}	0.231	0.528	0.905	0.134	0.609	0.859	0.944
multi-scale ^{M3}	0.208	0.446	0.791	0.119	0.656	0.890	0.961
NN blending ^{M2}	0.226	0.501	0.841	0.131	0.611	0.864	0.952
Mean blending ^{M2}	0.230	0.501	0.828	0.132	0.601	0.859	0.952
Frustum blending ^{M2}	0.229	0.499	0.826	0.131	0.604	0.861	0.953
Poisson blending ^{M2}	0.224	0.494	0.831	0.130	0.616	0.866	0.953

^{M2} Using MiDaS v2 [56]

^{M3} Using MiDaS v3 [55]




Figure 6. Estimated 360° depth maps at 2K resolution for indoor environments. Our results are closer to the ground-truth depth maps.




Figure 7. Estimated 360° disparity maps at 2048×1024 for outdoor environments [5]. Our results are more consistent geometrically.

We found in our experiments that blending disparity maps with the ‘frustum’ weights (see Figure 4) usually produces

results that are nearly as good as (see Table 3) but considerably faster than the Poisson blending of our complete method. This is a good compromise if speed is of essence. Concurrent to our work, Li et al. [40] use transformers for aligning and blending tangent depth maps based on predicted confidence.

Our proposed framework is the first to deal with high-resolution 360° images, and not limited to indoor scenes. Projecting the spherical input image onto a set of tangent images lets us overcome both the distortions of spherical projections and the resolution limits of deep monocular depth estimation methods. We proposed specially tailored optimisation techniques for global deformable multi-scale alignment and gradient-domain blending of the individual tangent disparity maps to overcome the discontinuous nature of tangent images. A major advantage of our approach is that we can leverage the high performance of MiDaS (or any future method) to generalise to new 360° datasets with higher accuracy and resolution than previous approaches. The resulting disparity maps at 2K resolution show a high level of geometric detail for both indoor and outdoor scenes.

Acknowledgements. This work was supported by the EPSRC CDT in Digital Entertainment (EP/L016540/1), an EPSRC-UKRI Innovation Fellowship (EP/S001050/1) and EPSRC grant CAMERA (EP/M023281/1, EP/T022523/1).

References

- [1] Sameer Agarwal, Keir Mierle, and Others. Ceres solver. <http://ceres-solver.org>, 2012.
- [2] Iro Armeni, Sasha Sax, Amir R. Zamir, and Silvio Savarese. Joint 2D-3D-semantic data for indoor scene understanding. [arXiv:1702.01105](https://arxiv.org/abs/1702.01105), 2017.
- [3] Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. MatryODShka: Real-time 6DoF video view synthesis using multi-sphere images. In *ECCV*, 2020.
- [4] Jiayang Bai, Shuichang Lai, Haoyu Qin, Jie Guo, and Yanwen Guo. GLPanoDepth: Global-to-local panoramic depth estimation. [arXiv:2202.02796](https://arxiv.org/abs/2202.02796), 2022.
- [5] Tobias Bertel, Mingze Yuan, Reuben Lindroos, and Christian Richardt. OmniPhotos: Casual 360° VR photography. *ACM Trans. Graph.*, 39(6):267:1–12, 2020.
- [6] 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. In *3DV*, pages 667–676, 2017.
- [7] Hong-Xiang Chen, Kunhong Li, Zhiheng Fu, Mengyi Liu, Zonghao Chen, and Yulan Guo. Distortion-aware monocular depth estimation for omnidirectional images. *IEEE Signal Processing Letters*, 28:334–338, 2021.
- [8] Hsien-Tzu Cheng, Chun-Hung Chao, Jin-Dong Dong, Hao-Kai Wen, Tyng-Luh Liu, and Min Sun. Cube padding for weakly-supervised saliency prediction in 360° videos. In *CVPR*, pages 1420–1429, 2018.
- [9] Taco S. Cohen, Mario Geiger, Jonas Koehler, and Max Welling. Spherical CNNs. In *ICLR*, 2018.
- [10] Taco S. Cohen, Maurice Weiler, Berkay Kicanaoglu, and Max Welling. Gauge equivariant convolutional networks and the icosahedral CNN. In *ICML*, 2019.
- [11] Benjamin Coors, Alexandru Paul Condurache, and Andreas Geiger. SphereNet: Learning spherical representations for detection and classification in omnidirectional images. In *ECCV*, pages 518–533, 2018.
- [12] James J. Cummings and Jeremy N. Bailenson. How immersive is enough? a meta-analysis of the effect of immersive technology on user presence. *Media Psychology*, 19(2):272–309, 2016.
- [13] Thiago Lopes Trugillo da Silveira and Claudio R. Jung. Dense 3D scene reconstruction from multiple spherical images for 3-DoF+ VR applications. In *IEEE VR*, pages 9–18, 2019.
- [14] 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 16×16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [15] Marc Eder, Pierre Moulon, and Li Guan. Pano popups: Indoor 3D reconstruction with a plane-aware network. In *3DV*, pages 76–84, 2019.
- [16] Marc Eder, Mykhailo Shvets, John Lim, and Jan-Michael Frahm. Tangent images for mitigating spherical distortion. In *CVPR*, 2020.
- [17] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, pages 2366–2374, 2014.
- [18] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning SO(3) equivariant representations with spherical CNNs. In *ECCV*, pages 52–68, 2018.
- [19] Clara Fernandez-Labrador, Jose M. Facil, Alejandro Perez-Yus, Cédric Demonceaux, Javier Civera, and Jose J. Guerrero. Corners for layout: End-to-end layout recovery from 360 images. *IEEE Robotics and Automation Letters*, 5(2):1255–1262, 2020.
- [20] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021.
- [21] Ravi Garg, Vijay Kumar B G, Gustavo Carneiro, and Ian Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *ECCV*, 2016.
- [22] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, pages 6602–6611, 2017.
- [23] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel Brostow. Digging into self-supervised monocular depth estimation. In *ICCV*, 2019.
- [24] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *ICCV*, pages 8977–8986, 2019.
- [25] Gaël Guennebaud, Benoît Jacob, and Others. Eigen v3. <https://eigen.tuxfamily.org>, 2010.
- [26] Peter Hedman and Johannes Kopf. Instant 3D photography. *ACM Trans. Graph.*, 37(4):101:1–12, 2018.
- [27] Derek Hoiem, Alexei A. Efros, and Martial Hebert. Automatic photo pop-up. *ACM Trans. Graph.*, 24(3):577–584, 2005.
- [28] Jingwei Huang, Zhili Chen, Duygu Ceylan, and Hailin Jin. 6-DOF VR videos with a single 360-camera. In *IEEE VR*, pages 37–44, 2017.
- [29] Sunghoon Im, Hyowon Ha, François Rameau, Hae-Gon Jeon, Gyeongmin Choe, and In So Kweon. All-around depth from small motion with a spherical panoramic camera. In *ECCV*, 2016.
- [30] 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.
- [31] Lei Jin, Yanyu Xu, Jia Zheng, Junfei Zhang, Rui Tang, Shugong Xu, Jingyi Yu, and Shenghua Gao. Geometric structure based and regularized depth estimation from 360 indoor imagery. In *CVPR*, pages 886–895, 2020.
- [32] Kevin Karsch, Ce Liu, and Sing Bing Kang. Depth transfer: Depth extraction from video using non-parametric sampling. *TPAMI*, 36(11):2144–2158, 2014.
- [33] Johannes Kopf, Kevin Matzen, Suhil Alsison, Ocean Quigley, Francis Ge, Yangming Chong, Josh Patterson, Jan-Michael Frahm, Shu Wu, Matthew Yu, Peizhao Zhang, Zijian He, Peter Vajda, Ayush Saraf, and Michael Cohen. One shot 3D photography. *ACM Trans. Graph.*, 39(4):76:1–13, 2020.

- [34] George Alex Koulieris, Kaan Akşit, Michael Stengel, Rafal K. Mantiuk, Katerina Mania, and Christian Richardt. Near-eye display and tracking technologies for virtual and augmented reality. *Comput. Graph. Forum*, 38(2):493–519, 2019.
- [35] Po Kong Lai, Shuang Xie, Jochen Lang, and Robert Laganière. Real-time panoramic depth maps from omni-directional stereo images for 6 DoF videos in virtual reality. In *IEEE VR*, pages 405–412, 2019.
- [36] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3DV*, 2016.
- [37] Yeonkun Lee, Jaeseok Jeong, Jongseob Yun, Wonjune Cho, and Kuk-Jin Yoon. SpherePHD: Applying CNNs on 360° images with non-euclidean spherical PolyHeDron representation. *TPAMI*, 44(2):834–847, 2022.
- [38] Junxuan Li, Hongdong Li, and Yasuyuki Matsushita. Lighting, reflectance and geometry estimation from 360° panoramic stereo. In *CVPR*, 2021.
- [39] Yuyan Li, Zhixin Yan, Ye Duan, and Liu Ren. PanoDepth: A two-stage approach for monocular omnidirectional depth estimation. In *3DV*, 2021.
- [40] Yuyan Li, Yuliang Guo, Zhixin Yan, Xinyu Huang, Ye Duan, and Liu Ren. OmniFusion: 360 monocular depth estimation via geometry-aware fusion. In *CVPR*, 2022.
- [41] Zhengqi Li and Noah Snavely. MegaDepth: Learning single-view depth prediction from internet photos. In *CVPR*, pages 2041–2050, 2018.
- [42] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T. Freeman. MannequinChallenge: Learning the depths of moving people by watching frozen people. *TPAMI*, 43(12):4229–4241, 2021.
- [43] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021.
- [44] Miaomiao Liu, Mathieu Salzmann, and Xuming He. Discrete-continuous depth estimation from a single image. In *CVPR*, 2014.
- [45] Thibault Louis, Jocelyne Troccaz, Amélie Rochet-Capellan, and François Bérard. Is it real? measuring the effect of resolution, latency, frame rate and jitter on the presence of virtual entities. In *International Conference on Interactive Surfaces and Spaces (ISS)*, pages 5–16, 2019.
- [46] Chenxu Luo, Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, Ram Nevatia, and Alan Yuille. Every pixel counts ++: Joint learning of geometry and motion with 3D holistic understanding. *TPAMI*, 42(10):2624–2641, 2020.
- [47] Yue Luo, Jimmy Ren, Mude Lin, Jiahao Pang, Wenxiu Sun, Hongsheng Li, and Liang Lin. Single view stereo matching. In *CVPR*, 2018.
- [48] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3D geometric constraints. In *CVPR*, pages 5667–5675, 2018.
- [49] S. Mahdi H. Miangoleh, Sebastian Dille, Long Mai, Sylvain Paris, and Yağız Aksoy. Boosting monocular depth estimation models to high-resolution via content-adaptive multi-resolution merging. In *CVPR*, 2021.
- [50] Pierre Moulon, Pascal Monasse, Romuald Perrot, and Renaud Marlet. OpenMVG: Open multiple view geometry. In *International Workshop on Reproducible Research in Pattern Recognition*, pages 60–74, 2016.
- [51] Grégoire Payen de La Garanderie, Amir Atapour Abarghouei, and Toby P. Breckon. Eliminating the blind spot: Adapting 3D object detection and monocular depth estimation to 360° panoramic imagery. In *ECCV*, pages 789–807, 2018.
- [52] Giovanni Pintore, Marco Agus, Eva Almansa, Jens Schneider, and Enrico Gobbetti. SliceNet: Deep dense depth estimation from a single indoor panorama using a slice-based representation. In *CVPR*, pages 11531–11540, 2021.
- [53] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, 2003.
- [54] Michaël Ramamonjisoa, Michael Firman, Jamie Watson, Vincent Lepetit, and Daniyar Turmukhambetov. Single image depth estimation using wavelet decomposition. In *CVPR*, 2021.
- [55] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, pages 12179–12188, 2021.
- [56] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *TPAMI*, 2021.
- [57] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J. Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *CVPR*, 2018.
- [58] Ashutosh Saxena, Min Sun, and Andrew Y. Ng. Make3D: Learning 3D scene structure from a single still image. *TPAMI*, 31(5):824–840, 2009.
- [59] Ana Serrano, Incheol Kim, Zhili Chen, Stephen DiVerdi, Diego Gutierrez, Aaron Hertzmann, and Belen Masia. Motion parallax for 360° RGBD video. *TVCG*, 25(5):1817–1827, 2019.
- [60] Jianping Shi, Xin Tao, Li Xu, and Jiaya Jia. Break Ames room illusion: Depth from general single images. *ACM Trans. Graph.*, 34(6):225:1–11, 2015.
- [61] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3D photography using context-aware layered depth inpainting. In *CVPR*, 2020.
- [62] Pratul P. Srinivasan, Rahul Garg, Neal Wadhwa, Ren Ng, and Jonathan T. Barron. Aperture supervision for monocular depth estimation. In *CVPR*, 2018.
- [63] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberley Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. [arXiv:1906.05797](https://arxiv.org/abs/1906.05797), 2019.
- [64] Yu-Chuan Su and Kristen Grauman. Learning spherical convolution for fast features from 360° imagery. In *NIPS*, 2017.

- [65] Yu-Chuan Su and Kristen Grauman. Kernel transformer networks for compact spherical convolution. In *CVPR*, pages 9442–9451, 2019.
- [66] Shinya Sumikura, Mikiya Shibuya, and Ken Sakurada. OpenVSLAM: a versatile visual SLAM framework. In *Proceedings of the International Conference on Multimedia*, 2019.
- [67] Cheng Sun, Min Sun, and Hwann-Tzong Chen. HoHoNet: 360 indoor holistic understanding with latent horizontal features. In *CVPR*, pages 2573–2582, 2021.
- [68] Richard Szeliski. Image alignment and stitching: a tutorial. *Foundations and Trends in Computer Graphics and Vision*, 2(1):1–104, 2006.
- [69] Keisuke Tateno, Nassir Navab, and Federico Tombari. Distortion-aware convolutional filters for dense prediction in panoramic images. In *ECCV*, pages 732–750, 2018.
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.
- [71] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *CVPR*, pages 2022–2030, 2018.
- [72] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web stereo video supervision for depth prediction from dynamic scenes. In *3DV*, 2019.
- [73] 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, 2018.
- [74] 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.
- [75] Fu-En Wang, Yu-Hsuan Yeh, Min Sun, Wei-Chen Chiu, and Yi-Hsuan Tsai. LED²-net: Monocular 360 layout estimation via differentiable depth rendering. In *CVPR*, 2021.
- [76] Ning-Hsu Wang, Bolivar Solarte, Yi-Hsuan Tsai, Wei-Chen Chiu, and Min Sun. 360SD-net: 360° stereo depth estimation with learnable cost volume. In *ICRA*, pages 582–588, 2020.
- [77] Changhee Won, Hochang Seok, Zhaopeng Cui, Marc Pollefeys, and Jongwoo Lim. OmniSLAM: Omnidirectional localization and dense mapping for wide-baseline multi-camera systems. In *ICRA*, pages 559–566, 2020.
- [78] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, and Zhenbo Luo. Monocular relative depth perception with web stereo data supervision. In *CVPR*, pages 311–320, 2018.
- [79] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3D scene shape from a single image. In *CVPR*, 2021.
- [80] Mingze Yuan and Christian Richardt. 360° optical flow using tangent images. In *BMVC*, 2021.
- [81] Wei Zeng, Sezer Karaoglu, and Theo Gevers. Joint 3D layout and depth prediction from a single indoor panorama image. In *ECCV*, 2020.
- [82] 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 *CVPR*, 2018.
- [83] Chao Zhang, Stephan Liwicki, William Smith, and Roberto Cipolla. Orientation-aware semantic segmentation on icosahedron spheres. In *ICCV*, pages 3533–3541, 2019.
- [84] Qiang Zhao, Chen Zhu, Feng Dai, Yike Ma, Guoqing Jin, and Yongdong Zhang. Distortion-aware CNNs for spherical images. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1198–1204, 2018.
- [85] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.
- [86] Chuanqing Zhuang, Zhengda Lu, Yiqun Wang, Jun Xiao, and Ying Wang. ACDNet: Adaptively combined dilated convolution for monocular panorama depth estimation. In *AAAI*, 2022.
- [87] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, and Petros Daras. OmniDepth: Dense depth estimation for indoors spherical panoramas. In *ECCV*, pages 448–465, 2018.
- [88] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, Federico Alvarez, and Petros Daras. Spherical view synthesis for self-supervised 360° depth estimation. In *3DV*, pages 690–699, 2019.