

Towards depth foundation models: Recent trends in vision-based depth estimation

Zhen Xu^{1,*}, Hongyu Zhou^{1,*}, Sida Peng¹, Haotong Lin¹, Haoyu Guo², Jiahao Shao¹, Peishan Yang¹, Qinglin Yang¹, Sheng Miao¹, Xingyi He¹, Yifan Wang¹, Yue Wang¹, Ruizhen Hu³, Yiyi Liao¹, Xiaowei Zhou¹, and Hujun Bao¹ (✉)

© The Author(s) 2026.

Abstract Depth estimation is a fundamental task in 3D computer vision, crucial for applications such as 3D reconstruction, free-viewpoint rendering, robotics, autonomous driving, and AR/VR technologies. Traditional methods relying on hardware sensors like LiDAR are often limited by their high costs, low resolution, and sensitivity to the environment, limiting their applicability to real-world scenarios. Recent advances in vision-based methods offer a promising alternative, yet they face challenges in generalization and stability due to either the low capacity of model architectures or reliance on domain-specific and small-scale datasets. The emergence of scaling laws and foundation models in other domains has inspired the development of “depth foundation models”: deep neural networks trained on large datasets with strong zero-shot generalization capabilities. This paper surveys the evolution of deep learning architectures and paradigms for depth estimation across monocular, stereo, multi-view, and monocular video settings. We explore the

potential of these models to address existing challenges and we also provide a comprehensive overview of large-scale datasets that can facilitate their development. By identifying key architectures and training strategies, we aim to highlight the path towards robust depth foundation models, offering insights for future research and applications.

Keywords depth estimation; foundation models; 3D computer vision

1 Introduction

Depth estimation is a cornerstone in the field of 3D computer vision. This task has been a focal point for researchers due to its critical role in various applications such as 3D reconstruction, 3D generative models, robotics, autonomous driving, and AR/VR technologies. However, algorithms often struggle to achieve high-quality, consistent depth recovery akin to human perception, which leverages prior knowledge of the scene and the world. Traditional depth recovery methods typically rely on active sensing hardware, including commercially available LiDAR sensors, time-of-flight (ToF) sensors, and ultrasonic probes. These sensors estimate depth by measuring the time taken for photons or sound waves to travel back and forth. Despite their accuracy, the high cost of these sensors limits their widespread application. Additionally, active sensors often suffer from low resolution and significant noise. For instance, the LiDAR sensor on an iPhone can only achieve a reconstruction resolution within a limited range and struggles with high precision for very close or distant objects. Moreover, these sensors are sensitive to

* Zhen Xu and Hongyu Zhou contributed equally to this work.

1 Zhejiang University, Hangzhou, China. E-mail: Z. Xu, zhenx@zju.edu.cn; H. Zhou, hongy_zhou@zju.edu.cn; S. Peng, pengsidea@zju.edu.cn; H. Lin, haotongl@outlook.com; J. Shao, jhshao@zju.edu.cn; P. Yang, yangps0306@gmail.com; Q. Yang, yangqinglin@zju.edu.cn; S. Miao, 12231099@zju.edu.cn; X. He, hexingyi8@gmail.com; Yifan Wang, accwyf@gmail.com; Yue Wang, ywang24@zju.edu.cn; Y. Liao, yiyi.liao@zju.edu.cn; X. Zhou, xwzhou@zju.edu.cn; H. Bao, bao@cad.zju.edu.cn (✉).

2 Shanghai AI Lab, Shanghai, China. E-mail: guohaoyu@pjlab.org.cn.

3 Shenzhen University, Shenzhen, China. E-mail: ruizhen.hu@szu.edu.cn.

Manuscript received: 2025-07-10; accepted: 2025-10-01



清华大学出版社
Tsinghua University Press

Available on
IEEE Xplore®

environmental lighting conditions, making them less effective outdoors in bright lighting.

Recently, there has been a growing interest in vision-based depth estimation methods that avoid active depth-sensing hardware, instead relying on readily available cameras commonly found in everyday devices. Unlike active sensor-based approaches, vision methods are cost-effective, offer an unrestricted depth range, are less affected by environmental conditions, and provide high resolution. For example, a standard iPhone camera can easily capture 4K resolution RGB information. However, existing vision-based depth estimation algorithms still face numerous challenges. Monocular depth estimation, in particular, is highly ill-posed, and standard deep learning algorithms struggle to achieve precise results. To introduce constraints and reduce ill-posedness, researchers often explore depth estimation algorithms with multiple camera inputs or scenarios with more extensive scene observations, such as stereo, multi-view, or video-based depth estimation. Nevertheless, these methods are often trained on small-scale synthetic data, leading to instability across spatial and temporal domains, poor generalization to different scenes and input types, and difficulties in overcoming the domain gap between synthetic and real-world data.

With the validation and rise of scaling laws in natural language processing, text-based image generation, and video generation models, the concept of foundation models has emerged. Foundation models are deep neural networks trained on vast amounts of data, exhibiting emergent zero-shot generalization capabilities in other domains. To achieve such capabilities, researchers focus on the scale and variation of input training data, leveraging large-scale models from other domains and ingeniously constructing self-supervision architectures. We define scalable depth estimation models and architectures capable of absorbing massive data as *depth foundation models*. For depth estimation tasks, including monocular, stereo, multi-view, and monocular video depth estimation, corresponding foundation models have the potential to address the aforementioned generalization issues and provide key solutions to long-standing computer vision tasks.

This paper aims to survey the evolution towards depth foundation models and paradigms for depth estimation across monocular, stereo, multi-view, and monocular video settings.

- We explore the development of deep learning model architectures and learning paradigms for each task and identify key paradigms with foundational capability or potential.
- To aid in the development of such depth foundation models, we also provide comprehensive surveys of large-scale datasets in each respective subfield.
- We also list the current key challenges faced by foundational architectures for each task to provide insight into future work.

2 Survey scope

2.1 Overview

This paper primarily concentrates on depth estimation methods that leverage deep learning, with a particular emphasis on foundation models that utilize large-scale architectures and extensive datasets. We begin by defining depth foundation models and then outline the depth estimation tasks that will be addressed in the following sections.

2.2 Definition of depth foundation models

We provide a brief overview of the development of foundation models in language models and self-supervised vision-based tasks to facilitate the understanding of the depth of foundation models. The field of language models has experienced explosive growth with the establishment of foundation models in recent years. This progress stems from the ability of these models to learn universal language and patterns from massive datasets, enabling them to generalize powerfully to various downstream tasks. Moreover, development of scalable architectures and self-supervision tasks has also enabled the emergence of vision foundation models, which are trained on a vast amount of image or video data to facilitate 2D and 3D perception tasks.

Convolutional neural networks and long short-term memory networks [1] played a main role in the early stages of language models, with limited network and data scales. The concept of word embeddings [2] and the introduction of the self-attention mechanism [3] allowed the model to process all words in a sequence simultaneously, vastly improving parallel computation efficiency and the ability to capture long-range dependencies. The original transformer model had a relatively small number of parameters, but

its architecture laid the groundwork for subsequent large-scale models. BERTs [4] and GPTs [5] can be considered as the beginning of foundation models of the large language model (LLM) type. Proposed by Google, BERT is a bidirectional pre-trained model based on the transformer architecture, enabling better understanding of the polysemy of words in a sentence. Bert was trained on the Toronto BookCorpus (800 million words) and the English Wikipedia (2.5 billion words). BERT-Base has 110 million parameters, and BERT-Large has 340 million parameters. GPT, proposed by OpenAI, is a unidirectional generative pre-trained model based on the transformer architecture. GPT models learn language patterns by predicting the next word, excelling in text generation tasks. GPT-3 was trained on a dataset over 45 TB in size, and has 175 billion parameters.

Vision foundation models have recently emerged as powerful tools for a wide range of visual perception tasks, closely following the development of scalable

model architectures in language models. These models, such as DINO [6–8], MAE [9], and SAM [10, 11], are typically trained on massive image or video datasets using self-supervised or weakly supervised objectives. By leveraging scalable architectures (e.g., vision transformers) and large-scale pretraining, vision foundation models learn rich and generalizable visual representations that can be transferred to downstream tasks, including classification, detection, segmentation, and even 3D perception. Notably, recent diffusion-based vision models [12, 13] have also demonstrated strong capabilities in generative and dense prediction tasks. These advances in vision foundation models provide the architectural and data-driven basis for the development of depth foundation models, enabling improved generalization and scalability for depth estimation in diverse domains.

The development of depth estimation models is illustrated in Fig. 1. Considering the scale of *language* foundation models, we define a *depth* foundation




Fig. 1 Scaling trends in model capacity and data volume for depth estimation. Each point represents a published method, positioned by its approximate model size (horizontal axis, $\log(\text{number of parameters})$) and dataset size (vertical axis, $\log(\text{number of training images})$), and colored by publication year. Early models (lightest shading, before 2020) relied on sub-million-parameter networks trained on only thousands of images, yielding limited generalization. In 2020 and 2021 (light blue shading), methods grew to tens of millions of parameters and hundreds of thousands of image datasets. In 2022 and 2023 (gray shading), methods continued to increase in size and scalable training, with around one billion parameters and datasets ranging from several million to ten million images. The most recent models (darkest shading, 2024 and 2025) further scale up to billions of parameters and utilize over ten million images, with strong generalization ability, and demonstrate the potential evolution toward depth foundation models.

model as one trained on a large dataset (over 10 million images) and employs models with a substantial number of parameters (over 1 billion). Additionally, depth foundation models should exhibit strong generalizability to multiple data domains.

2.3 Depth estimation tasks

This survey covers several tasks, including monocular depth estimation, stereo depth estimation, multi-view depth estimation, and monocular video depth estimation using foundation models. Let $\mathbf{I} = \{I_{k,t}, k=1, \dots, \mathcal{K}, t=1, \dots, \mathcal{T}\}$ represent a collection of RGB images, where $\mathcal{K}$ denotes the number of cameras and $\mathcal{T}$ is the number of timestamps for the frames. In the case of monocular depth estimation, the input consists of a single image $I_{1,1}$. For stereo depth estimation, the input comprises a pair of images $\{I_{1,1}, I_{2,1}\}$. In multi-view depth estimation, the input is a set of images captured at the same timestamp but varying spatial locations, represented as $\{I_{k,1}, k=1, \dots, \mathcal{K}\}$. For monocular video depth estimation, the input consists of a sequence of images captured by a monocular camera at different timestamps, represented as $\{I_{1,t}, t=1, \dots, \mathcal{T}\}$. The scope of our survey excludes the task of multi-view video depth estimation, which can be represented as the most general form of input: $\{I_{k,t}, k=1, \dots, \mathcal{K}, t=1, \dots, \mathcal{T}\}$. This is due to the fact that foundation models for this task have not yet been thoroughly explored.

Moreover, we additionally define the task of metric depth estimation as one that requires the model to not only output a relative depth map, but also predict its real-world scale and offset. For the tasks of stereo, multi-view, and monocular video depth estimation, depth estimation requires predicting depth maps that, when unprojected to the world coordinate system, align across different views. When the real-world metric-scale camera parameters are known, these three tasks automatically become metric depth estimation tasks. Thus, we only distinguish relative and metric depth estimation in the monocular depth estimation task.

For each task, we begin by reviewing the background and evolution of deep learning models specific to the task. We then delve into the development of foundation models. Prominent examples of foundation models include transformer-based models and diffusion models. Furthermore,

we also discuss the large-scale datasets used for training these foundation models, encompassing both synthetic and real-world datasets, which enable the models to generalize effectively to diverse scenes. Finally, we address problems faced by existing depth foundation models.

3 Overview of depth estimation

3.1 Paradigms

In monocular image depth estimation, models have progressed from direct depth regression to affine-invariant depth, depth classification, and canonical camera depth. This progression facilitates more accurate depth map predictions using just a single input image. In stereo image depth estimation, by utilizing the principles of stereo geometry, models can concentrate on matching corresponding pixels in image pairs. This has driven the evolution of paradigms from cost-volume methods to the attention mechanism, iterative optimizers, and ultimately to scalable training approaches. In multi-view image depth estimation, similar to stereo paradigms, the incorporation of multi-view information has facilitated an evolution from patch-match stereo to cost volume methods, followed by a coarse-to-fine strategy and the implementation of token attention mechanisms. In monocular image depth estimation, by incorporating an additional dimension of timestamps, models leverage temporal correlation and test-time optimization paradigms to establish the relationship between temporal and spatial information. Additionally, scaling up training enhances depth estimation performance and guides models toward becoming depth foundation models.

3.2 Datasets


We summarize the datasets commonly used in depth estimation tasks in Table 1. These datasets can be categorized into two classes: real-world-captured and synthetic. Each dataset may be applicable to multiple tasks, and we specify the tasks associated with each dataset in the table. Furthermore, we present details of the scenes within the datasets, including whether they provide metric information, whether they include dynamic scenes, the number of scenes and frames, the image resolution, and whether the annotations are dense, based on the proportion of pixels with corresponding depth values.

Table 1 Datasets. Scenes in datasets marked  are object-centric. Similarly,  and  refer to indoor and outdoor scenes, respectively. Depth annotations can be classified as dense or sparse, according to whether most pixels have corresponding depth values. †: HM3D is a digital twin dataset created from real-world data, with the number of frames and resolutions being user-defined

Domain	Task				Dataset name	Dyn.	Metric	Scene#	Frame#	Resolution	Anno.	License
	Mono	Stereo	MV	Video								
Real-world	✓				A2D2 [14] 	Dyn.	Yes	3	394k	1920 × 1208	Sparse	CC BY-ND 4.0
	✓	✓			ArogoVerse2 [15, 16] 	Dyn.	Yes	1000	2.14M	2048 × 1550	Sparse	CC BY-NC-SA 4.0
	✓	✓			CityScapes [17] 	Dyn.	Yes	2975	89k	2048 × 1024	Sparse	Non-commercial
	✓				DDAD [18] 	Dyn.	Yes	200	98k	1936 × 1216	Sparse	CC BY-NC-SA 4.0
	✓	✓			DIML [19]  	Static	Yes	200+	2M	1334 × 756	Sparse	Non-commercial
	✓				DIODE [20]  	Static	Yes	30	27k	1024 × 768	Sparse	MIT
	✓	✓			DSEC [21] 	Dyn.	Yes	53	26k	1440 × 1080	Sparse	CC BY-NC-SA 4.0
	✓				HM3D [22] 	Static	Yes	1000	†	†	Dense	MIT
	✓				iBims-1 [23] 	Static	Yes	20	54	640 × 480	Sparse	CC BY 4.0
	✓				Lyft [24] 	Dyn.	Yes	366	158k	1224 × 1024	Sparse	CC0 1.0
	✓				Mapillary PSD [25] 	Dyn.	Yes	50k	750k	1920 × 1080	Sparse	CC BY-NC-SA
	✓				NuScenes [26] 	Dyn.	Yes	1k	40k	1600 × 900	Sparse	Non-commercial
	✓				NYU [27] 	Static	Yes	464	435k	640 × 480	Dense	GPL
	✓				Pandaset [28] 	Dyn.	Yes	103	8k	1920 × 1080	Sparse	CC0
	✓				Taskonomy [29] 	Static	Yes	600	4.5M	512 × 512	Dense	MIT
	✓	✓			KITTI [30] 	Dyn.	Yes	22	41k	1242 × 375	Sparse	CC BY-NC-SA 3.0
	✓	✓			DrivingStereo [31] 	Dyn.	Yes	184	180k	1762 × 800	Sparse	MIT
		✓			InStereo2K [32] 	Static	Yes	50	2k	1080 × 860	Dense	Non-commercial
		✓			Argoverse [33] 	Dyn.	Yes	113	6624	2056 × 2464	Sparse	CC BY-NC-SA 4.0
		✓			ETH3D [34]  	Static	Yes	25	1k	6233 × 4146	Dense	CC BY-NC-SA 4.0
			✓		Waymo [35] 	Dyn.	Yes	1150	160k	1920 × 1280	Sparse	Non-commercial
			✓		UASOL [36] 	Static	Yes	676	160.9k	2280 × 1282	Dense	CC BY-NC-SA 3.0
			✓		DTU [37] 	Static	Yes	124	42.5k	1600 × 1200	Dense	MIT
			✓		BlendedMVS [38]   	Static	No	113	17k	2048 × 1536	Dense	CC BY 4.0
	✓			✓	WildRGBD [39] 	Static	Yes	23k	6M	480 × 640	Dense	MIT
	✓			✓	MVImgNet [40] 	Static	No	220k	6.8M	1080 × 1920	Dense	CC BY-NC 4.0
	✓			✓	ARKitScenes [41] 	Static	Yes	5047	450M	256 × 192	Dense	Non-commercial
	✓			✓	ARKitScenes-HD [41] 	Static	Yes	5047	450M	1920 × 1440	Dense	Non-commercial
	✓			✓	Matterport3D [42] 	Static	Yes	90	194.4k	1280 × 1024	Dense	MIT
	✓	✓			Replica [43] 	Static	Yes	18	36k	1200 × 680	Dense	Non-commercial
	✓	✓		✓	Dyn. Replica [44] 	Dyn.	Yes	484	145k	1280 × 720	Dense	CC BY-NC 4.0
	✓			✓	ScanNet [45] 	Static	Yes	1513	2.5M	640 × 480	Dense	Non-commercial
	✓			✓	ScanNet++ [46] 	Static	Yes	1858	3.7M+	1920 × 1440	Dense	Non-commercial
Synthetic	✓			✓	Sintel [47]  	Dyn.	No	10	1k	1024 × 436	Dense	CC BY 3.0
		✓			SceneFlow [48]   	Dyn.	No	9	40k	960 × 540	Dense	CC BY 4.0
		✓			CREStereo [49] 	Static	No	0	103k	1920 × 1080	Dense	Apache-2.0
		✓			FallingThings [50]  	Dyn.	No	3	62k	960 × 540	Dense	CC BY-NC-SA 4.0
		✓			FSD [51]   	Dyn.	No	12	1M	1280 × 720	Dense	Non-commercial
		✓			UnrealStereo4K [52]  	Static	No	8	7720	3840 × 2160	Dense	MIT
		✓		✓	Spring [53] 	Dyn.	No	47	6k	1920 × 1080	Dense	CC BY 4.0
		✓		✓	TartanAir [54]  	Static	Yes	163	1M	640 × 480	Dense	CC BY 4.0
		✓		✓	VirtualKITTI2 [55] 	Dyn.	Yes	5	25k	1242 × 375	Dense	CC BY-NC-SA 4.0
			✓	✓	MatrixCity [56] 	Static	Yes	3k	519k	1000 × 1000	Dense	Apache-2.0
			✓	✓	MVS-Synth [57] 	Dyn.	No	120	12k	1920 × 1080	Dense	Non-commercial
				✓	3D Ken Burns [58] 	Static	No	32	536k	512 × 512	Dense	CC BY-NC-SA 4.0
	✓			✓	OmniObject3D [59] 	Static	Yes	6k	600k	800 × 800	Dense	CC BY 4.0
	✓	✓		✓	IRS [60] 	Static	Yes	70	100k	960 × 540	Dense	Apache 2.0
			✓	✓	PointOdyssey [61]  	Dyn.	No	131	200k	540 × 960	Dense	MIT
			✓	✓	BEDLAM [62]   	Dyn.	Yes	10,450	380k	1280 × 720	Dense	Non-commercial

Table 2 Key methods in depth estimation. Model and data sizes are approximate. Benchmark results are collected from corresponding papers. Sections marked with different background colors (light gray or medium gray) indicate different evaluation protocols used by the papers

Task	Method	Year	Model type	Key paradigm	Model size	Data size	Benchmark	
							NYUv2 [27] (Rel.) ↓	KITTI [30, 63] (Rel.) ↓
Monocular	Multi-Scale Net [64]	2014	CNN	Direct regression	~70M	~150k	—	0.190
	LRC [65]	2016	CNN	Direct regression	~30M	~50k	—	0.114
	DORN [66]	2018	CNN	Depth classification	~50M	~200k	0.115	0.072
	PAD-Net [67]	2018	CNN	Multiple tasks	~25M	~3k	0.120	—
	MiDaS [68]	2019	CNN	Affine invariant	~90M	~2M	0.111	0.236
	LeReS [69]	2020	CNN	Canonical camera	~90M	~350k	0.090	0.149
	DPT [70]	2021	ViT	Scaling law	~350M	~1.5M	0.098	0.100
	Marigold [71]	2024	Diffusion	Scaling law	~800M	~75k	0.055	0.099
	DA [72]	2024	ViT	Scaling law	~600M	~63.5M	0.0382	0.0504
	DAV2 [73]	2024	ViT	Scaling law	~2B	~62.5M	0.0416	0.0677
	MoGE [74]	2024	ViT	Scaling law	~350M	~4M	0.0292	0.0394
	UniDepth [75]	2024	ViT	Metric depth	~350M	~4M	0.0340/0.116	0.0350/0.0469
	UniDepthV2 [76]	2025	ViT	Metric depth	~350M	~4M	0.0296/0.106	0.0385/0.0858
	MoGE2 [77]	2025	ViT	Metric depth	~350M	~4M	0.0289/0.073	0.0375/0.1810
Stereo	KITTI 2015 [63] (D1) ↓							
	GCNet [78]	2017	CNN	Cost volume	~4M	~35k	2.87%	
	STTR [79]	2020	ViT	Attention	~3M	~22k	2.01%	
	LEAStereo [80]	2020	CNN	Attention	~2M	~20k	1.65%	
	RAFT-Stereo [81]	2021	CNN+GRU	Iterative optim.	~10M	~35k	1.96%	
	Mono2Stereo [82]	2024	CNN	Scaling law	~5M	~600k	1.58%	
Multi-view	DTU [83] (Overall Err.) ↓							
	Gipuma [84]	2015	Handcrafted	Handcrafted algo.	~0	0	0.283	
	MVSNet [85]	2018	CNN	Cost volume	~4M	~27k	0.462	
	DUST3R [86]	2023	ViT	Scaling law	~400M	~17M	1.741	
	MVSFormer++ [87]	2024	ViT	Scaling law	~40M	~40k	0.281	
	MASt3R [88]	2024	ViT	Scalable law	~400M	~20M	0.374	
	VGGt [89]	2025	ViT	Scaling law	~1B	~30M	0.382	
Video	ScanNet [45] (Rel.) ↓ KITTI [30, 63] (Rel.) ↓							
	CVD [90]	2020	CNN	Test-time optim.	~25M	~120k	0.073	—
	MAMo [91]	2023	ViT	Temporal corr.	~200M	~60k	—	0.049
	NVDS [92]	2023	ViT	Scaling law	~100M	~2.2M	0.187	0.253
	ChronoDepth [93]	2024	Diffusion	Scaling law	~2B	~240k	0.159	0.151
	DepthCrafter [94]	2024	Diffusion	Scaling law	~2B	~25M	0.125	0.110
	Buffer Anytime [95]	2024	ViT+Diffusion	Scaling law	~300M	~26M	0.123	0.119

**Fig. 2** Evolution of key ideas in monocular image depth estimation. From early direct regression and classification methods, through affine-invariant and canonical camera depth estimation, monocular depth estimation models have shown increasingly stronger generalization capabilities, paving the way for the emergence of depth foundation models.

4 Monocular image depth estimation

4.1 Overview

Monocular depth estimation aims to predict per-pixel depth distances for a scene from a single RGB image, establishing a geometric mapping from 2D images to 3D scenes. The core challenges lie in scale ambiguity and the lack of geometric information inherent in monocular vision. Monocular depth estimation requires no additional hardware and relies solely on image content to infer scene structure, making it valuable for various applications such as 3D reconstruction, video editing, autonomous driving, and augmented reality. In recent years, with the advancement of deep learning techniques, monocular depth estimation has gradually shifted from traditional geometric methods to data-driven end-to-end learning paradigms, evolving toward foundation models with strong generalization capabilities, high accuracy, and robustness.

4.2 Evolution of model architectures

The evolution of monocular depth estimation has been driven by progressive innovations in network architectures, depth representations, and learning paradigms. Early approaches predominantly employed handcrafted image processing pipelines [96–99], with seminal works like CNN [64] establishing the deep learning foundation. Subsequent refinements incorporated U-Net [100] and ResNet [101] residual blocks to enhance spatial continuity. The advent of dense prediction transformers [70] marked a paradigm shift: DPT [70] introduced global attention mechanisms through patch-wise sequence modeling, addressing CNNs’ inherent locality constraints. The latest frontier involves diffusion models [102, 103], exemplified by Marigold [71], which leverage conditional denoising frameworks to transfer generative priors into depth estimation, significantly improving geometric consistency and cross-domain generalization.




Fig. 3 Representative pipelines for monocular image depth estimation. The pink components denote operations without learnable parameters or with fixed parameters, while the blue components indicate operations with optimizable parameters. Vision transformer-based approaches, leveraging their lightweight architectures, enable real-time monocular depth estimation. However, due to the presence of convolution operations in their architectures, they may lose detailed features. Diffusion model-based methods treat RGB images as conditional inputs, effectively preserving fine-grained details. However, their denoising processes impose computational costs, making it challenging to achieve real-time performance.




Fig. 4 Evolution of the stereo image depth estimation paradigm. Paradigms have transitioned from cost-volume methods to attention mechanisms and iterative optimization techniques to effectively match the features of stereo images. The incorporation of monocular and diffusion priors facilitates large-scale training, paving the way for foundation models.

4.3 Evolution of method paradigms

Parallel advances emerged in depth representation and learning strategies. Early methods [64, 100, 101] focused on absolute depth prediction but faced scale ambiguity challenges, prompting innovations like scale-invariant loss [68, 104] for relative depth estimation. Another advancement came through bin classification approaches [66, 105–109], which discretize the continuous depth space into multiple bins. This reformulates depth estimation as a per-pixel classification problem, effectively handling non-uniform depth distributions and capturing uncertainty in predictions, particularly at depth discontinuities. Recent breakthroughs [69, 75, 110–112] integrate camera parameters to resolve metric ambiguity while maintaining scale awareness. The data paradigm evolved through large-scale multi-view stereo (MVS) datasets like MegaDepth [113], enabling unprecedented generalization. Supervision methods expanded beyond fully labeled training: self-supervised approaches exploited photometric consistency in stereo sequences [114], while Depth Anything [72] demonstrated the effectiveness of pseudo-label distillation for knowledge transfer. Multi-task frameworks [67, 115–117] further enhanced robustness through joint depth-flow-pose estimation with geometric constraints, complemented by geometry-aware losses like those based on virtual normal that explicitly enforce surface regularity. This multifaceted progression underscores how architectural innovation, representation learning, and supervision paradigms have collectively advanced monocular depth estimation towards human-level scene understanding.

4.4 Evolution towards foundation models

Models for monocular depth estimation are becoming larger and more data-intensive, evolving towards foundation models. In terms of *model size*, advanced foundation models primarily utilize two key methods: vision transformer (ViT) architectures [118] and diffusion-based generative models [13]. The dense prediction transformer (DPT) provides the main architecture for modern depth estimation [70], with 343 million parameters. Unlike older fully convolutional networks, DPT’s ViT backbone keeps high-resolution features and a global view throughout, leading to more detailed and consistent depth estimates. Pre-trained diffusion models are used as strong depth estimators to improve the working of monocular depth estimation [71, 119–121], thanks to their existing visual knowledge, which ranges from 200 million to 1 billion parameters. For example, Marigold [71] modifies text-to-image latent diffusion models (like Stable Diffusion v2 [12]) to predict depth from images. In terms of *data scale*, powerful hardware and better cameras have led to many high-quality depth datasets, typically ranging from thousands to millions of images; further details can be found in Table 1. The current trend is to train models on large datasets to make them more adaptable. When training for large-scale depth estimation, models typically use loss functions that ignore scale and shift differences in the data [68]. This means models learn to predict relative depth, making relative depth estimation a standard task in computer vision. However, metric depth estimation, which provides absolute distance measurements, is crucial for real-world applications. Currently, there are two main ways to get metric depth: (i) methods like those in Refs. [74, 75, 77, 110–

112, 122] combine camera intrinsic estimates with relative depth predictions to get metric outputs through geometric calculations, and (ii) directly learning metric depth: approaches such as those in Refs. [123–125] train models directly on data that includes scale information, allowing them to predict metric depth without extra steps.

4.5 Outstanding problems

Recent methods have made notable progress, yet critical challenges persist in four areas.

- *Depth accuracy* enhancement remains the foremost pursuit, aiming to develop a “visual LiDAR” system in which RGB cameras rival dedicated depth sensors. Current methods exhibit 4%–5% relative error on standard benchmarks [75], escalating to 20%–50% in challenging scenarios [126], while hardware sensors consistently achieve better than 1% accuracy.
- *Absolute scale recovery* has been improved through works like Metric3D [111], but results are fragile under complex illumination and in texture-deficient conditions, necessitating more robust geometric priors.
- *Data efficiency* is hindered by compounded limitations: pseudo-label noise restricts supervision quality while synthetic-to-real domain gaps constrain model generalization, demanding innovative low-cost high-precision annotation paradigms.
- *Multi-task generalization* presents an open research frontier, as current approaches struggle to unify depth estimation with complementary tasks like semantic segmentation and surface normal prediction within a single foundational model architecture. These interconnected challenges collectively underscore the need for fundamental breakthroughs in geometric understanding, data utilization, and cross-task knowledge integration to bridge the performance gap between learning-based methods and physical sensing systems.

5 Stereo image depth estimation

5.1 Overview

Stereo depth estimation aims to estimate the per-pixel depth for a scene given the relative pose of a pair of stereo cameras and a pair of RGB observations as input. Compared to monocular

depth estimation, stereo depth estimation can utilize disparity priors to estimate depth through epipolar feature matching and triangulation, which theoretically yields better results than monocular depth estimation. Additionally, since stereo depth estimation incorporates the known relative pose of the cameras, the estimated depth inherently includes scale information, addressing the scale ambiguity present in monocular depth estimation. The core challenge of stereo depth estimation lies in accurately establishing the correspondence between pixels in the left and right images. The traditional matching process can be easily affected by factors such as changes in lighting, occlusion, repetitive textures, and weak textures.

With the rapid development of deep learning technology, in stereo depth estimation, neural networks are gradually replacing traditional epipolar feature matching methods. The main neural network architecture paradigms include the CNN-based cost-volume paradigm, the transformer-based attention paradigm, and the RNN-based iterative optimization paradigm.

5.2 Evolution of model architectures

CNN-based models provide one of the most classic and widely used technical approaches to stereo depth estimation [52, 78, 80, 127–134]. Early methods, such as PSMNet [128], used simple convolutional layers to extract features, while subsequent works adopted more complex architectures like ResNet [135], DenseNet [136], and NAS search techniques [80] to extract more robust features. Since 2020, with the success of transformers [3] in the field of language models and the advances of vision transformers [118] in computer vision, the attention paradigm has been considered [79, 137–142]. To reduce computational complexity, Ref. [137] introduced local attention and sparse attention mechanisms. Inspired by RNNs [143] and the application of the RAFT network architecture to optical flow estimation, Refs. [81, 144, 145] constructed a correlation pyramid and used an RNN model to iteratively optimize the disparity map. Ref. [145] introduced learnable update strategies that dynamically adjust the update direction through attention mechanisms. Refs. [51, 146–149] utilized monocular foundation models as priors and diffusion models as inpainting tools for stereo generation, respectively, advancing stereo depth estimation through large-scale training.






Fig. 5 Representative pipelines of stereo image depth estimation. The first architecture allows for scalable training, which leverages all available datasets along with pseudo stereo pairs synthesized from a monocular dataset, to train a foundation model. The second architecture migrates knowledge from a monocular foundation model to a stereo model, making it possible to achieve a stereo foundation model from relatively small-scale training datasets.

5.3 Evolution of method paradigms

The CNN-based cost volume paradigm [52, 78, 80, 128–134] constructs a cost volume using 2D or 3D CNNs to represent a matching cost between left and right images, and then performs cost aggregation as a post-processing step to ultimately regress the disparity map. The attention paradigm [79, 137–142, 145] models the epipolar matching problem as a sequence-to-sequence task, utilizing self-attention and cross-attention mechanisms to capture long-range dependencies between pixels. Within a single image, the self-attention mechanism models the relationships between pixels to capture contextual information, while the cross-attention mechanism matches corresponding pixels between left and right images. Specifically, each pixel in the left image interacts with all pixels in the right image to compute attention weights. The iterative optimization paradigm [81, 144, 145] dynamically updates the cost volume or feature maps based on the current disparity map to capture more accurate

matching information. It can balance performance and speed through early stopping, making it suitable for time-sensitive applications. Additionally, this paradigm exhibits strong robustness to initial errors. With significant advances in simulation technology, generative techniques, and monocular depth estimation methods in recent years, foundation models [51, 54, 147–151] in the field of stereo depth estimation are beginning to emerge.

5.4 Evolution towards foundation models

Foundation models are emerging as the new paradigm for stereo depth estimation, leading to an increase in data intensity. However, in contrast to foundation models used for monocular depth estimation, the *model size* for stereo depth estimation has not seen a significant increase, ranging from 3.5 million to 11 million parameters [78–81]. Instead, advances in monocular depth estimation foundation models allow stereo tasks to leverage monocular priors to enhance depth estimation. There are currently two main approaches to utilizing these priors. One approach

[51, 146] involves injecting features from monocular depth models into the cost volume. Another approach [147, 148] applies stereo metrics to scale monocular depth estimates and then uses refinement networks to provide better depth estimation results. There are two methods to enable *large data scale training* for stereo depth estimation: (i) creating high-fidelity virtual scenes using advanced simulation and rendering technologies, as done by FSD [51] (1 million image pairs) and TartanAir [54] (1 million image pairs), and (ii) synthesizing stereo data from monocular data. Learning stereo from single images [151] (597 thousand image pairs) and Stereo anything [150] (30 million image pairs) generate virtual stereo pairs by using estimated scene depth from monocular images [68, 72], allowing pixels to be warped to pseudo-stereo viewpoints and inpainting to fill in gaps. Recent advances in diffusion models have led to stereoGen [149] (35 thousand image pairs) using stable diffusion as an inpainting tool.

5.5 Outstanding problems

Research into foundation models for stereo depth estimation is still in its early stages, and we believe there are several key issues worth exploring in future:

- *Limited data.* Generated stereo data from monocular faces challenges such as insufficient accuracy in monocular depth estimation and difficulties in warping to fill holes. Additionally, there remains a domain gap between synthetic data and the real world, and synthetic datasets still lack diversity.
- *Lack of end-to-end training.* Current methods that leverage monocular priors treat monocular foundation models as cues, lacking end-to-end training of large model parameters. The model parameter counts are relatively small, and there is a lack of foundation model designs tailored for stereo tasks.
- *Limited utilization of available datasets.* There is insufficient application of cross-domain datasets. Aside from Stereo Anything [150], most approaches typically utilize only one or two datasets, failing to fully leverage existing stereo datasets.
- *Under-utilization of the diffusion architecture and priors.* A diffusion foundation model for stereo depth estimation, similar to the Marigold model [71], has not been explored.

6 Multi-view image depth estimation

6.1 Overview

Multi-view image depth estimation (more commonly known as multi-view stereo, MVS) aims to recover the 3D structure of a scene from multiple images captured from different viewpoints, leveraging known or estimated camera poses. This task is fundamental to 3D reconstruction and underpins applications such as robotics, AR/VR, and autonomous navigation. Compared to monocular and stereo settings, the multi-view scenario provides richer geometric cues and enables more accurate and complete scene understanding, but also introduces new challenges in terms of view selection, feature aggregation, and computational efficiency. It is worth noting that estimated depth maps are typically already aligned with the camera parameters from multiple viewpoints; thus, we do not distinguish between relative and metric depth estimation in this section. Over the years, the field has evolved from early heuristic and patch-based matching methods, through cost volume-based deep learning approaches with 2D and 3D CNNs, to recent advances utilizing transformers and diffusion models. These developments have not only improved the robustness and accuracy of depth estimation, for example, in sparse-view cases, but also paved the way for large-scale pretraining and the emergence of foundation models with strong generalization capabilities. In this section, we review the key technical paradigms, architectural innovations, and open challenges in multi-view image depth estimation, highlighting the trajectory towards foundation models and the remaining obstacles for real-world deployment.

6.2 Evolution of model architectures

The model architectures for multi-view depth estimation have evolved from early heuristic matching to state-of-the-art transformer and diffusion-based models. Early approaches [152, 153] typically assumed that camera parameters were already known, obtained from SfM [154–156], SLAM [157, 158], or robotic arm control [159, 160], and rely on traditional dense pixel matching techniques [161] for depth estimation. This heuristic matching laid the groundwork for future developments. With the advent of deep learning, researchers began to






Fig. 6 Evolution of the multi-view image depth estimation paradigm. From early heuristic matching, through 2D CNNs and 3D CNNs, to advanced frameworks utilizing transformers and diffusion models, multi-view depth estimation models have shown increasing robustness and accuracy.

employ 2D CNNs [162, 163] to process the cost volume, which has low computational cost and enables real-time scene reconstruction. At the same time, alternative approaches utilized 3D CNNs [57, 85, 145, 164–170] to achieve more accurate depth estimation. Recent trends have started to explore the use of transformers [87–89, 171–179] or diffusion models [180], which leverage global self-attention mechanisms or diffusion-based strategies to capture richer feature representations and enhance inference quality.

6.3 Evolution of method paradigms

The method paradigms for multi-view depth estimation have evolved from early PatchMatch-based matching strategies to modern token attention mechanisms. Traditional methods [152, 153] relied on the PatchMatch algorithm [161] for pixel-level matching—a simple yet influential approach that set the stage for subsequent innovations. With the integration of deep learning, constructing a cost volume by back-projecting multi-view features using camera parameters became mainstream [57, 85, 162, 164, 165], allowing for more accurate depth regression by computing feature correlations. Given the high memory consumption of cost volumes, modern approaches [145, 166–170] adopt a coarse-to-fine strategy, starting with a low resolution and a large number of depth hypotheses, and progressively refining the cost volume with predictions from previous stages to achieve a balance between high resolution and high accuracy. Looking ahead, emerging research is investigating the incorporation of token attention [86, 88, 89, 173, 175–178, 181–186] mechanisms in the depth estimation process, aiming

to better capture both local and global contexts and offering promising new directions for multi-view depth estimation.

6.4 Evolution towards foundation models


In terms of *model size*, recent advances have been driven by the emergence of transformer-based foundation models. MVSTR [171] was the first to introduce the transformer architecture to multi-view depth estimation, employing a global-context transformer and a 3D-geometry transformer for intra-view global feature extraction and inter-view information interaction. MVSFormer [172] proposed use of a pre-trained vision transformer (ViT) to enhance multi-view stereo tasks, leveraging priors learned from large-scale datasets. MVSFormer++ [87] further employed a pre-trained DINOv2 model with 1.1 billion parameters and introduced a distinct attention mechanisms tailored for the feature encoder and cost volume regularization. Both PF-LRM [187] and DUST3R [86] utilized ViT to directly predict point maps, with DUST3R employing the CroCo model featuring a complete encoder-decoder transformer architecture. In terms of *data scale*, the evolution shows a clear trend from small-scale datasets to large-scale training. Early methods like MVSNet [85] used only 27k images from the DTU dataset, while MVSFormer++ [87] used 40k images from the DTU and BlendedMVS datasets. The breakthrough came with sparse-view methods: PF-LRM [187] utilized 1M objects from Objaverse and MVImgNet datasets, DUST3R [86] was trained on 17M image pairs from eight datasets, including Habitat, MegaDepth, and Waymo, GRM [173] used 40M objects from Objaverse, and VGGT [89] employed 30M images from

- *Depth estimation of objects with complex materials.* Reflective or transparent scenes pose significant challenges for geometry estimation due to their complex optical properties. Incorporating learned priors into depth estimation presents a promising approach for accurately capturing the geometry of complex materials.

7 Monocular video depth estimation

Video depth estimation aims to estimate per-frame depth from a given monocular video while ensuring temporal consistency across the sequence. Its primary challenge over monocular depth estimation lies in maintaining consistency over time. In contrast to multi-view stereo (MVS) [87, 173] and other multi-view depth estimation methods [171, 172], it needs to handle dynamic scenes, which imposes further challenges. Since video depth estimation requires both temporally consistent and accurate depth predictions,

- *Sparse view reconstruction.* In real-world applications, capturing a scene with dense and complete views may be infeasible due to constraints on reachability. Therefore, the ability to reconstruct complete scene geometry from partial observations by leveraging prior knowledge represents an important direction for future research.
- *Fine-grained depth estimation.* Current feed-forward methods [87, 171–173] can efficiently predict multi-view image depth in a single forward pass. However, accurately capturing fine-grained geometry remains challenging, as it demands high-precision geometric prediction capabilities from neural networks.






Fig. 8 Evolution of monocular video depth estimation. From RNN-based temporal modeling (2018) to CNNs with test-time optimization (CVD, 2020), transformer-based scaling (NVDS, 2023), and video diffusion for enhanced stability and generalization.

and its technical paradigm integrates elements from both monocular and multi-view depth estimation, we consider it the ultimate problem that a depth foundation model should address.

7.2 Evolution of model architectures

With the advancement of deep learning technologies and network architectures, various network structures have been applied in recent years to improve the temporal consistency of depth estimation. With the rise and widespread adoption of RNNs in language models, in 2019, Ref. [188] proposed using the LSTM mechanism to integrate temporal information, thereby enhancing the stability of video depth estimation. Meanwhile, in image tasks, CNN architectures (including U-Net and ResNet) have shown excellent performance. In 2020, Ref. [90] introduced test-time optimization to single-frame depth estimation methods based on CNN architectures, amplifying the capabilities of traditional CNNs through optimization and alignment. Subsequently, attention mechanisms gained significant attention. Refs. [91, 175–177] applied attention mechanisms to integrate temporal features, maintaining memory through attention mechanisms rather than LSTMs, thereby achieving more accurate and consistent depth estimation results. Later, with the remarkable success of diffusion models in the field of image generation and the impressive results of video diffusion models, in 2024, Refs. [93, 94] proposed use of video diffusion models to enhance consistency, achieving unprecedented stability and predictive performance.

7.3 Evolution of method paradigms

In recent years, innovative technical paradigms have been proposed to enhance the stability of video depth

estimation. In 2019, Ref. [188] introduced a memory mechanism to enable networks to integrate multi-frame information, implicitly encoding content from other frames to assist in depth estimation for the current frame. Subsequently, several works explored similar memory paradigms, such as MAMo [91]. In 2020, Ref. [90] proposed using test-time optimization to post-process predicted depth results. Since this approach better leverages prior techniques like SfM and SLAM while also performing bundle adjustment (BA) on depth, it often yields superior results to purely generalized methods. Later, with the rise of scaling laws, interest in generalized methods was reignited, and attention shifted to the generation and utilization of large-scale training data. In 2023, Ref. [92] introduced a network-based post-processing method for monocular depth estimation results, finding a middle ground between direct video depth prediction and test-time training (TTT). It also introduced the first representative large-scale video depth in the wild (VDW) dataset. Video depth estimation has since begun to evolve toward the development of foundation models.

7.4 Evolution towards foundation models

In terms of *model size*, recent work has focused on using scalable network architectures to model video depth. Buffer Anytime [95] proposed using a temporal ViT architecture with 343 million parameters to integrate temporal information and generate a large amount of pseudo ground truth (GT) based on single-image priors. Specifically, the input images are first split into patches and converted into tokens recognizable by ViT. Then, self-attention is applied to these tokens. To fuse temporal information, cross-attention is performed on the single-frame




Fig. 9 Representative pipelines of monocular video depth estimation. The vision transformer (ViT) paradigm processes inputs split into patches through self-attention and temporal attention mechanisms to integrate spatial and temporal depth cues. The video diffusion paradigm leverages denoising diffusion models, using concatenated depth and image features, optionally enriched with CLIP embeddings, to generate consistent video depth estimates. These scalable architectures enhance generalization and enable zero-shot depth estimation across diverse datasets.

attention results, reprocessing the features to obtain tokens integrated with temporal information. Finally, the tokens are decoded and patches recombined to produce the final depth estimation results. Another technical paradigm [72, 93, 94] proposed using video diffusion models with 200 million to 1 billion parameters to aid in recovering video depth. Specifically, these methods use the input video as a condition in the video diffusion denoising process. During denoising, the noisy depth is concatenated with the corresponding frame's RGB color and, optionally, with CLIP embeddings as input to the denoising UNet. In terms of *data scale*, the lack of sufficient video depth data has always been a critical issue [189]. In 2023, NVDS [92] constructed the first large-scale, diverse video stereo depth dataset (14.2k videos, 2.24M frames) by using a video stereo matching method to generate pseudo labels. In 2024, methods like Refs. [72, 93, 94] addressed the lack of video data by leveraging video priors from video generation models. Following NVDS [92],

DepthCrafter [94] annotated 200k videos using a video stereo matching method. Depth Any Video [72] created 40k videos, totaling 6M frames, using a game engine. Video Depth Anything [190] proposed using an image teacher to provide pseudo labels to compensate for the lack of video depth ground truth (GT) labels.

7.5 Outstanding problems

- *Geometric inconsistency.* When camera motion is present in a monocular video, the ability to estimate consistent depth for the same statistical scene observed across multiple frames is crucial. To achieve this, jointly modeling camera motion and depth estimation is a promising approach and future direction for enhancing geometric consistency.
- *Temporal inconsistency.* In the presence of dynamic objects, the lack of multi-view observation in monocular video makes it particularly challenging to estimate their geometry. The

method needs to learn strong priors to predict temporally consistent depths for dynamic objects across multiple frames.

- *Monocular video depth training data.* Learning strong dynamic priors needs extensive and diverse training data. However, due to the presence of dynamic objects, collecting real-world ground-truth monocular video training data with accurate ground-truth depths often requires additional depth sensors, limiting the scale of collection. Scaling up monocular video training data, either using real-world Internet unlabeled video with a self-supervised training strategy, or simulating and rendering realistic synthetic data, is a direction worth exploring.

8 Applications

8.1 Overview

Depth foundation models are expected to see a wide range of practical applications across computer vision, robotics, and graphics. By providing accurate, generalizable, and scalable depth estimation from monocular images, stereo pairs, multi-view images, or video sequences, these models could serve as a critical component for downstream tasks that require 3D scene understanding. In this section, we highlight several representative application domains in which depth foundation models would make a significant impact, including 3D reconstruction, novel view synthesis, video world modeling, robotics, and autonomous driving.

8.2 3D reconstruction

For multi-view 3D reconstruction tasks, once the depths for each view have been estimated, they can be directly fused to obtain the reconstructed scene geometry, which is usually represented as a point cloud or triangle mesh. MVSNet [85] and CasMVSNet [167] adopt Fusibile [84], which first converts the depth map of each view into a point cloud, then filters out points with poor consistency by projecting them into other views and checking, and finally fuses the point clouds from all views to obtain the reconstruction of the entire scene. SimpleRecon [191] and Murre [180] use TSDF fusion [192], which converts depth maps into sparse truncated SDF grids, averages them across multiple views, and finally extracts the surface

using marching cubes to obtain a triangle mesh. Direct fusion methods are relatively efficient but require high-quality depth maps. NeuRIS [193] and MonoSDF [194] use relative depths predicted from a monocular depth estimation method as a supervision signal, constraining the SDF field through a specially designed depth loss, and thereby enhancing the quality of the 3D reconstruction.

8.3 Novel view synthesis

In recent years, NeRF-based [195] and 3D Gaussian splatting-based [196] methods have made significant progress in the task of novel view synthesis. High-quality depth estimates can serve as strong priors for these models, enhancing their performance and accelerating convergence. Specifically, the depth map can be utilized as sampling guidance for NeRF-based methods [194, 197, 198] and as a coarse initialization for 3D Gaussian splatting-based methods [199–201]. Additionally, the dense depth map can also be employed to learn a dense SDF field during the training process, allowing for the alignment of the geometry of NeRFs or 3D Gaussians with the SDF field. This alignment can improve the geometry of the reconstructed results and facilitate the synthesis of high-quality novel views, particularly for perspectives that are distant from the training views. Some works [202–206] depend on LiDAR point clouds or RGB-D images as geometric priors. However, both LiDAR point clouds and RGB-D images can be costly to acquire, requiring additional sensors. High-quality depth estimates can serve as a substitute for these methods, unleashing the potential of novel view synthesis techniques.

8.4 Video world models

As the popularity of diffusion models grows, video diffusion models have been proposed to generalize the image synthesis pipeline to video generation. With the success of SORA [207] and other video foundation models [208, 209], several works [210–212] have explored video models' capabilities as world models. Having a foundational model for depth estimation, preferably on videos, would significantly bridge the gap between image-only generation models' understanding of 3D space and motion. As the ability to predict depth would indicate, the model would at least possess the capability of distinguishing the size and placement pattern of everyday objects, scenes,

and people. Having depth cues for video generation models could potentially serve as a breaking point for further boosting current world models' ability to understand everyday scenes and might even stimulate the underlying generalization ability even further, leading to more spatially and temporally consistent generation and prediction results.

8.5 Robotics and autonomous driving

In robotics and autonomous driving, accurate and reliable depth perception plays a pivotal role in tasks such as navigation, obstacle detection, and collision avoidance. Traditional solutions often rely on LiDAR or stereo camera systems, both of which come with increased hardware costs and complexity. Depth foundation models learned from large-scale datasets have the potential to deliver high-quality depth estimates from a single camera, making them particularly attractive for cost-sensitive real-world applications. Recent methods [125] demonstrate that monocular depth estimation can be integrated into robotic perception pipelines, serving to complement or even substitute for more expensive sensors. For instance, some works [110, 111] adopt monocular depth estimation to enhance SLAM or visual odometry frameworks, showing improvements in localization and mapping under challenging lighting or weather conditions. Similarly, in autonomous driving, depth estimates can facilitate large-scale 3D reconstructions for building realistic simulation environments, supporting algorithm development and testing [202]. Furthermore, by leveraging depth priors learned from diverse scenes, these models exhibit promising generalization capabilities, potentially enabling robust domain adaptation to varied environments, from urban streets to off-road terrains, thus paving the way for more versatile and scalable robotic and self-driving solutions.

9 Future work

9.1 Overview

The previous sections discuss several fundamental problems to be solved before we reach a general-purpose depth foundation model. These revolve around data and consistency.

9.2 Data

For all of the discussed depth estimation tasks:

monocular image, stereo image, multi-view image, and monocular video depth estimation, the lack of accurate, large-scale, high-quality, and highly variable data is currently the main obstacle to constructing and training a depth foundation model. Due to the unique nature of the depth estimation task, current approaches to acquiring data usually fall into two categories: using depth sensors, and synthetic rendering. The main depth sensors used are LiDAR and ultrasonic devices. However, the acquired ground truth depth maps are usually incomplete or noisy due to the sensitive nature of the depth sensing devices. In synthetic data generation, several attempts have been made to curate high-quality, hand-crafted, large static or dynamic 3D scenes by artists. However, these data are naturally limited to a small scale due to the amount of work required. Future work should focus on either utilizing self-supervision techniques to better transfer the knowledge of vast image and video data to the task of depth estimation, or developing a better approach for simulation and generation, providing artist-quality synthetic rendering and depth pairs to boost generalization ability.

9.3 Consistency

Consistency should be both spatial and temporal. For the task of monocular image depth estimation, current methods typically fall short when merging depth estimation results from different timestamps and viewpoints for the same scene. Although a vast amount of work has investigated the issue of temporal consistency throughout the target video, video depth estimation still fails to produce accurate and consistent results when given multiple viewpoints of the same 3D scene or trying to unproject and merge the prediction results [125]. Notably, multi-view video reconstruction methods [213–215] have proved the existence of dynamic and multi-view inductive bias of the 4D world by providing accurate reconstructions from only image-based optimization objects. Future work should focus on exploring the intrinsic 3D or dynamic inductive bias present in the dynamic 3D world, to further mitigate the problem of spatial and temporal inconsistency.

10 Conclusions

Since 2022, advances in foundation models within the natural language processing domain, along with the



emergence of scaling laws, have led to a significant increase in the development of foundation models in the field of computer vision. In recent years, numerous foundation models have been introduced for depth estimation tasks, and new models continue to emerge at a rapid pace, challenging practitioners to stay up-to-date with the latest developments.

In this timely paper, we present a comprehensive survey of foundation models for depth estimation tasks, covering their background, development, and latest advances. We also address outstanding problems faced by existing depth foundation models and their downstream applications. We hope this paper serves as a valuable guide for practitioners and researchers interested in depth estimation foundation models.

Finally, numerous challenges and opportunities remain in the realm of depth foundation models. We believe that as foundation models continue to evolve and depth estimation tasks advance, we will witness an increasing number of sophisticated and practical applications in the future.

Author contributions

Zhen Xu and Hongyu Zhou were responsible for the overall writing of the manuscript. Sida Peng, Yiyi Liao, Yue Wang, Ruizhen Hu, Xiaowei Zhou, and Hujun Bao provided critical supervision and guidance throughout the project, shaping its framework, refining technical discussions, and ensuring clarity and coherence, providing valuable oversight and feedback during drafting and revision. The remaining co-authors supported the work by evaluating key publications and charting the evolution timeline of depth estimation architectures across the monocular, stereo, multi-view, and monocular video depth estimation tasks to deliver a thorough survey.

Acknowledgements

This research was supported by Zhejiang Provincial Natural Science Foundation of China (LD25F030001), and the Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

References

- [1] Graves, A. Long short-term memory. In: *Supervised Sequence Labelling with Recurrent Neural Networks. Studies in Computational Intelligence, Vol. 385*. Springer Berlin Heidelberg, 37–45, 2012.
- [2] Almeida, F.; Xexéo, G. Word embeddings: A survey. *arXiv preprint arXiv:1901.09069*, 2019.
- [3] Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; Polosukhin, I. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [4] Devlin, J.; Chang, M. W.; Lee, K.; Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Vol. 1 (Long and Short Papers)*, 4171–4186, 2019.
- [5] Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [6] Caron, M.; Touvron, H.; Misra, I.; Jegou, H.; Mairal, J.; Bojanowski, P.; Joulin, A. Emerging properties in self-supervised vision transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9630–9640, 2022.
- [7] Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. DINOv2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- [8] Siméoni, O.; Vo, H. V.; Seitzer, M.; Baldassarre, F.; Oquab, M.; Jose, C.; Khalidov, V.; Szafraniec, M.; Yi, S.; Ramamonjisoa, M.; et al. DINOv3. *arXiv preprint arXiv:2508.10104*, 2025.
- [9] He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; Girshick, R. Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15979–15988, 2022.
- [10] Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W. Y.; et al. Segment anything. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3992–4003, 2024.
- [11] Ravi, N.; Gabeur, V.; Hu, Y. T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. SAM 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.



- [12] Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; Ommer, B. High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10674–10685, 2022.
- [13] Ho, J.; Jain, A.; Abbeel, P. Denoising diffusion probabilistic models. *arXiv preprint* arXiv:2006.11239, 2020.
- [14] Geyer, J.; Kassahun, Y.; Mahmudi, M.; Ricou, X.; Durgesh, R.; Chung, A. S.; Hauswald, L.; Pham, V. H.; Mühlegg, M.; Dorn, S.; et al. A2D2: Audi autonomous driving dataset. *arXiv preprint* arXiv:2004.06320, 2020.
- [15] Wilson, B.; Qi, W.; Agarwal, T.; Lambert, J.; Singh, J.; Khandelwal, S.; Pan, B.; Kumar, R.; Hartnett, A.; Pontes, J. K.; et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. *arXiv preprint* arXiv:2301.00493, 2023.
- [16] Lambert, J.; Hays, J. Trust, but verify: Cross-modality fusion for HD map change detection. *arXiv preprint* arXiv:2212.07312, 2021.
- [17] Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; Schiele, B. The cityscapes dataset for semantic urban scene understanding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3213–3223, 2016.
- [18] Guizilini, V.; Ambrus, R.; Pillai, S.; Raventos, A.; Gaidon, A. 3D packing for self-supervised monocular depth estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2482–2491, 2020.
- [19] Cho, J.; Min, D.; Kim, Y.; Sohn, K. DIML/CVL RGB-D dataset: 2M RGB-D images of natural indoor and outdoor scenes. *arXiv preprint* arXiv:2110.11590, 2021.
- [20] Vasiljevic, I.; Kolkin, N.; Zhang, S.; Luo, R.; Wang, H.; Dai, F. Z.; Daniele, A. F.; Mostajabi, M.; Basart, S.; Walter, M. R.; et al. DIODE: A dense indoor and outdoor DEpth dataset. *arXiv preprint* arXiv:1908.00463, 2019.
- [21] Gehrig, M.; Aarents, W.; Gehrig, D.; Scaramuzza, D. DSEC: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters* Vol. 6, No. 3, 4947–4954, 2021.
- [22] Ramakrishnan, S. K.; Gokaslan, A.; Wijmans, E.; Maksymets, O.; Clegg, A.; Turner, J.; Undersander, E.; Galuba, W.; Westbury, A.; Chang, A. X.; et al. Habitat-matterport 3D dataset (HM3D): 1000 large-scale 3D environments for embodied AI. *arXiv preprint* arXiv:2109.08238, 2021.
- [23] Koch, T.; Liebel, L.; Fraundorfer, F.; Körner, M. Evaluation of CNN-based single-image depth estimation methods. In: *Computer Vision – ECCV 2018. Lecture Notes in Computer Science, Vol. 11131*. Leal-Taixé, L.; Roth, S. Eds. Springer Cham, 331–348, 2019.
- [24] Houston, J.; Zuidhof, G.; Bergamini, L.; Ye, Y.; Chen, L.; Jain, A.; Omari, S.; Iglovikov, V.; Ondruska, P. One thousand and one hours: Self-driving motion prediction dataset. *arXiv preprint* arXiv:2006.14480, 2020.
- [25] Antequera, M. L.; Gargallo, P.; Hofinger, M.; Bulò, S. R.; Kuang, Y.; Kotschieder, P. Mapillary planet-scale depth dataset. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science, Vol. 12347*. Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J.-M. Eds. Springer Cham, 589–604, 2020.
- [26] Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; Beijbom, O. nuScenes: A multimodal dataset for autonomous driving. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11618–11628, 2020.
- [27] Silberman, N.; Hoiem, D.; Kohli, P.; Fergus, R. Indoor segmentation and support inference from RGBD images. In: *Computer Vision – ECCV 2012. Lecture Notes in Computer Science, Vol. 7576*. Fitzgibbon, A.; Lazebnik, S.; Perona, P.; Sato, Y.; Schmid, C. Eds. Springer Berlin Heidelberg, 746–760, 2012.
- [28] Xiao, P.; Shao, Z.; Hao, S.; Zhang, Z.; Chai, X.; Jiao, J.; Li, Z.; Wu, J.; Sun, K.; Jiang, K.; et al. PandaSet: Advanced sensor suite dataset for autonomous driving. In: *Proceedings of the IEEE International Intelligent Transportation Systems Conference*, 3095–3101, 2021.
- [29] Zamir, A. R.; Sax, A.; Shen, W.; Guibas, L.; Malik, J.; Savarese, S. Taskonomy: Disentangling task transfer learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3712–3722, 2018.
- [30] Geiger, A.; Lenz, P.; Urtasun, R. Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3354–3361, 2012.
- [31] Yang, G.; Song, X.; Huang, C.; Deng, Z.; Shi, J.; Zhou, B. DrivingStereo: A large-scale dataset for stereo matching in autonomous driving scenarios. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 899–908, 2020.
- [32] Bao, W.; Wang, W.; Xu, Y.; Guo, Y.; Hong, S.;



- Zhang, X. InStereo2K: A large real dataset for stereo matching in indoor scenes. *Science China Information Sciences* Vol. 63, No. 11, Article No. 212101, 2020.
- [33] Chang, M. F.; Lambert, J.; Sangkloy, P.; Singh, J.; Bak, S.; Hartnett, A.; Wang, D.; Carr, P.; Lucey, S.; Ramanan, D.; et al. Argoverse: 3D tracking and forecasting with rich maps. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8740–8749, 2020.
- [34] Schöps, T.; Schönberger, J. L.; Galliani, S.; Sattler, T.; Schindler, K.; Pollefeys, M.; Geiger, A. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2538–2547, 2017.
- [35] Sun, P.; Kretzschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; et al. Scalability in perception for autonomous driving: Waymo open dataset. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2443–2451, 2020.
- [36] Bauer, Z.; Gomez-Donoso, F.; Cruz, E.; Orts-Escolano, S.; Cazorla, M. UASOL, a large-scale high-resolution outdoor stereo dataset. *Scientific Data* Vol. 6, Article No. 162, 2019.
- [37] Aanæs, H.; Jensen, R. R.; Vogiatzis, G.; Tola, E.; Dahl, A. B. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision* Vol. 120, No. 2, 153–168, 2016.
- [38] Yao, Y.; Luo, Z.; Li, S.; Zhang, J.; Ren, Y.; Zhou, L.; Fang, T.; Quan, L. BlendedMVS: A large-scale dataset for generalized multi-view stereo networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1787–1796, 2020.
- [39] Xia, H.; Fu, Y.; Liu, S.; Wang, X. RGBD objects in the wild: Scaling real-world 3D object learning from RGB-D videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22378–22389, 2024.
- [40] Yu, X.; Xu, M.; Zhang, Y.; Liu, H.; Ye, C.; Wu, Y.; Yan, Z.; Zhu, C.; Xiong, Z.; Liang, T.; et al. MVIImgNet: A large-scale dataset of multi-view images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9150–9161, 2023.
- [41] Baruch, G.; Chen, Z.; Dehghan, A.; Dimry, T.; Feigin, Y.; Fu, P.; Gebauer, T.; Joffe, B.; Kurz, D.; Schwartz, A.; et al. ARKitScenes: A diverse real-world dataset for 3D indoor scene understanding using mobile RGB-D data. *arXiv preprint* arXiv:2111.08897, 2021.
- [42] Chang, A.; Dai, A.; Funkhouser, T.; Halber, M.; Nießner, M.; Savva, M.; Song, S.; Zeng, A.; Zhang, Y. Matterport3D: Learning from RGB-D data in indoor environments. *arXiv preprint* arXiv:1709.06158, 2017.
- [43] Straub, J.; Whelan, T.; Ma, L.; Chen, Y.; Wijmans, E.; Green, S.; Engel, J. J.; Mur-Artal, R.; Ren, C.; Verma, S.; et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint* arXiv:1906.05797, 2019.
- [44] Karaev, N.; Rocco, I.; Graham, B.; Neverova, N.; Vedaldi, A.; Rupprecht, C. DynamicStereo: Consistent dynamic depth from stereo videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13229–13239, 2023.
- [45] Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; Nießner, M. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2432–2443, 2017.
- [46] Yeshwanth, C.; Liu, Y. C.; Nießner, M.; Dai, A. ScanNet++: A high-fidelity dataset of 3D indoor scenes. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12–22, 2024.
- [47] Butler, D. J.; Wulff, J.; Stanley, G. B.; Black, M. J. A naturalistic open source movie for optical flow evaluation. In: *Computer Vision – ECCV 2012. Lecture Notes in Computer Science, Vol. 7577*. Fitzgibbon, A.; Lazebnik, S.; Perona, P.; Sato, Y.; Schmid, C. Eds. Springer Berlin Heidelberg, 611–625, 2012.
- [48] Mayer, N.; Ilg, E.; Häusser, P.; Fischer, P.; Cremers, D.; Dosovitskiy, A.; Brox, T. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4040–4048, 2016.
- [49] Li, J.; Wang, P.; Xiong, P.; Cai, T.; Yan, Z.; Yang, L.; Liu, J.; Fan, H.; Liu, S. Practical stereo matching via cascaded recurrent network with adaptive correlation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16242–16251, 2022.
- [50] Tremblay, J.; To, T.; Birchfield, S. Falling things: A synthetic dataset for 3D object detection and pose estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2119–21193, 2018.
- [51] Wen, B.; Trepte, M.; Aribido, J.; Kautz, J.; Gallo, O.; Birchfield, S. FoundationStereo: Zero-shot stereo matching. *arXiv preprint* arXiv:2501.09898, 2025.



- [52] Tosi, F.; Liao, Y.; Schmitt, C.; Geiger, A. SMD-nets: Stereo mixture density networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8938–8948, 2021.
- [53] Mehl, L.; Schmalfluss, J.; Jahedi, A.; Nalivayko, Y.; Bruhn, A. Spring: A high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 4981–4991, 2023.
- [54] Wang, W.; Zhu, D.; Wang, X.; Hu, Y.; Qiu, Y.; Wang, C.; Hu, Y.; Kapoor, A.; Scherer, S. TartanAir: A dataset to push the limits of visual SLAM. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 4909–4916, 2020.
- [55] Cabon, Y.; Murray, N.; Humenberger, M. Virtual KITTI 2. *arXiv preprint* arXiv:2001.10773, 2020.
- [56] Li, Y.; Jiang, L.; Xu, L.; Xiangli, Y.; Wang, Z.; Lin, D.; Dai, B. MatrixCity: A large-scale city dataset for city-scale neural rendering and beyond. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 3182–3192, 2024.
- [57] Huang, P. H.; Matzen, K.; Kopf, J.; Ahuja, N.; Huang, J. B. DeepMVS: Learning multi-view stereopsis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2821–2830, 2018.
- [58] Niklaus, S.; Mai, L.; Yang, J.; Liu, F. 3D Ken Burns effect from a single image. *ACM Transactions on Graphics* Vol. 38, No. 6, Article No. 184, 2019.
- [59] Wu, T.; Zhang, J.; Fu, X.; Wang, Y.; Ren, J.; Pan, L.; Wu, W.; Yang, L.; Wang, J.; Qian, C.; et al. OmniObject3D: Large-vocabulary 3D object dataset for realistic perception, reconstruction and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 803–814, 2023.
- [60] Wang, Q.; Zheng, S.; Yan, Q.; Deng, F.; Zhao, K.; Chu, X. IRS: A large naturalistic indoor robotics stereo dataset to train deep models for disparity and surface normal estimation. In: Proceedings of the IEEE International Conference on Multimedia and Expo, 1–6, 2021.
- [61] Zheng, Y.; Harley, A. W.; Shen, B.; Wetzstein, G.; Guibas, L. J. PointOdyssey: A large-scale synthetic dataset for long-term point tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 19798–19808, 2024.
- [62] Black, M. J.; Patel, P.; Tesch, J.; Yang, J. BEDLAM: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8726–8737, 2023.
- [63] Menze, M.; Geiger, A. Object scene flow for autonomous vehicles. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 3061–3070, 2015.
- [64] Eigen, D.; Puhrsch, C.; Fergus, R. Depth map prediction from a single image using a multi-scale deep network. In: Proceedings of the 28th International Conference on Neural Information Processing Systems, Vol. 2, 2366–2374, 2014.
- [65] Godard, C.; Mac Aodha, O.; Brostow, G. J. Unsupervised monocular depth estimation with left-right consistency. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 6602–6611, 2017.
- [66] Fu, H.; Gong, M.; Wang, C.; Batmanghelich, K.; Tao, D. Deep ordinal regression network for monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2002–2011, 2018.
- [67] Xu, D.; Ouyang, W.; Wang, X.; Sebe, N. PAD-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 675–684, 2018.
- [68] Ranftl, R.; Lasinger, K.; Hafner, D.; Schindler, K.; Koltun, V. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 44, No. 3, 1623–1637, 2022.
- [69] Yin, W.; Zhang, J.; Wang, O.; Niklaus, S.; Mai, L.; Chen, S.; Shen, C. Learning to recover 3D scene shape from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 204–213, 2021.
- [70] Ranftl, R.; Bochkovskiy, A.; Koltun, V. Vision transformers for dense prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 12159–12168, 2022.
- [71] Ke, B.; Obukhov, A.; Huang, S.; Metzger, N.; Daut, R. C.; Schindler, K. Repurposing diffusion-based image generators for monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 9492–9502, 2024.
- [72] Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; Zhao, H. Depth anything: Unleashing the power of large-scale unlabeled data. In: Proceedings of the

- IEEE/CVF Conference on Computer Vision and Pattern Recognition, 10371–10381, 2024.
- [73] Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; Zhao, H. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024.
- [74] Wang, R.; Xu, S.; Dai, C.; Xiang, J.; Deng, Y.; Tong, X.; Yang, J. MoGe: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 5261–5271, 2025.
- [75] Piccinelli, L.; Yang, Y. H.; Sakaridis, C.; Segu, M.; Li, S.; Gool, L. V.; Yu, F. UniDepth: Universal monocular metric depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 10106–10116, 2024.
- [76] Piccinelli, L.; Sakaridis, C.; Yang, Y. H.; Segu, M.; Li, S.; Abbeloos, W.; Van Gool, L. UniDepthV2: Universal monocular metric depth estimation made simpler. *arXiv preprint arXiv:2502.20110*, 2025.
- [77] Wang, R.; Xu, S.; Dong, Y.; Deng, Y.; Xiang, J.; Lv, Z.; Sun, G.; Tong, X.; Yang, J. MoGe-2: Accurate monocular geometry with metric scale and sharp details. *arXiv preprint arXiv:2507.02546*, 2025.
- [78] Kendall, A.; Martirosyan, H.; Dasgupta, S.; Henry, P.; Kennedy, R.; Bachrach, A.; Bry, A. End-to-end learning of geometry and context for deep stereo regression. In: Proceedings of the IEEE International Conference on Computer Vision, 66–75, 2017.
- [79] Li, Z.; Liu, X.; Drenkow, N.; Ding, A.; Creighton, F. X.; Taylor, R. H.; Unberath, M. Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 6177–6186, 2022.
- [80] Cheng, X.; Zhong, Y.; Harandi, M.; Dai, Y. Hierarchical neural architecture search for deep stereo matching. In: Proceedings of the 34th International Conference on Neural Information Processing Systems, 22158–22169, 2020.
- [81] Lipson, L.; Teed, Z.; Deng, J. RAFT-stereo: Multilevel recurrent field transforms for stereo matching. In: Proceedings of the International Conference on 3D Vision, 218–227, 2021.
- [82] Wang, Y.; Liang, Y.; Li, H.; Fu, Y. Mono2Stereo: Monocular knowledge transfer for enhanced stereo matching. *arXiv preprint arXiv:2411.09151*, 2024.
- [83] Jensen, R.; Dahl, A.; Vogiatzis, G.; Tola, E.; Aanaes, H. Large scale multi-view stereopsis evaluation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 406–413, 2014.
- [84] Galliani, S.; Lasinger, K.; Schindler, K. Massively parallel multiview stereopsis by surface normal diffusion. In: Proceedings of the IEEE International Conference on Computer Vision, 873–881, 2016.
- [85] Yao, Y.; Luo, Z.; Li, S.; Fang, T.; Quan, L. MVSNet: Depth inference for unstructured multi-view stereo. In: *Computer Vision – ECCV 2018. Lecture Notes in Computer Science, Vol. 11212*. Ferrari, V.; Hebert, M.; Sminchisescu, C.; Weiss, Y. Eds. Springer Cham, 785–801, 2018.
- [86] Wang, S.; Leroy, V.; Cabon, Y.; Chidlovskii, B.; Revaud, J. DUSt3R: Geometric 3D vision made easy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 20697–20709, 2024.
- [87] Cao, C.; Ren, X.; Fu, Y. MVFormer++: Revealing the devil in Transformer’s details for multi-view stereo. *arXiv preprint arXiv:2401.11673*, 2024.
- [88] Leroy, V.; Cabon, Y.; Revaud, J. Grounding image matching in 3D with MAST3R. In: *Computer Vision – ECCV 2024. Lecture Notes in Computer Science, Vol. 15083*. Leonardis, A.; Ricci, E.; Roth, S.; Russakovsky, O.; Sattler, T.; Varol, G. Eds. Springer Cham, 71–91, 2025.
- [89] Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Ruppert, C.; Novotny, D. VGGT: Visual geometry grounded transformer. *arXiv preprint arXiv:2503.11651*, 2025.
- [90] Luo, X.; Huang, J. B.; Szeliski, R.; Matzen, K.; Kopf, J. Consistent video depth estimation. *ACM Transactions on Graphics Vol. 39, No. 4, Article No. 71*, 2020.
- [91] Yasarla, R.; Cai, H.; Jeong, J.; Shi, Y.; Garrepalli, R.; Porikli, F. MAMo: Leveraging memory and attention for monocular video depth estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 8720–8730, 2024.
- [92] Wang, Y.; Shi, M.; Li, J.; Huang, Z.; Cao, Z.; Zhang, J.; Xian, K.; Lin, G. Neural video depth stabilizer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 9432–9442, 2024.
- [93] Shao, J.; Yang, Y.; Zhou, H.; Zhang, Y.; Shen, Y.; Guizilini, V.; Wang, Y.; Poggi, M.; Liao, Y. Learning temporally consistent video depth from video diffusion priors. *arXiv preprint arXiv:2406.01493*, 2024.
- [94] Hu, W.; Gao, X.; Li, X.; Zhao, S.; Cun, X.; Zhang, Y.; Quan, L.; Shan, Y. DepthCrafter: Generating consistent long depth sequences for open-world videos. *arXiv preprint arXiv:2409.02095*, 2024.
- [95] Kuang, Z.; Zhang, T.; Zhang, K.; Tan, H.; Bi, S.; Hu, Y.; Xu, Z.; Hasan, M.; Wetzstein, G.; Luan, F. Buffer anytime: Zero-shot video depth and normal

- from image priors. *arXiv preprint arXiv:2411.17249*, 2024.
- [96] Saxena, A.; Chung, S. H.; Ng, A. Y. Learning depth from single monocular images. In: Proceedings of the 19th International Conference on Neural Information Processing Systems, 1161–1168, 2005.
- [97] Hoiem, D.; Efros, A. A.; Hebert, M. Recovering surface layout from an image. *International Journal of Computer Vision* Vol. 75, No. 1, 151–172, 2007.
- [98] Liu, C.; Yuen, J.; Torralba, A.; Sivic, J.; Freeman, W. T. SIFT flow: Dense correspondence across different scenes. In: *Computer Vision – ECCV 2008. Lecture Notes in Computer Science, Vol. 5304*. Forsyth, D.; Torr, P.; Zisserman, A. Eds. Springer Berlin Heidelberg, 28–42, 2008.
- [99] Saxena, A.; Sun, M.; Ng, A. Y. Make3D: Learning 3D scene structure from a single still image. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 31, No. 5, 824–840, 2009.
- [100] Eigen, D.; Fergus, R. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: Proceedings of the IEEE International Conference on Computer Vision, 2650–2658, 2016.
- [101] Laina, I.; Rupprecht, C.; Belagiannis, V.; Tombari, F.; Navab, N. Deeper depth prediction with fully convolutional residual networks. In: Proceedings of the 4th International Conference on 3D Vision, 239–248, 2016.
- [102] Ji, Y.; Chen, Z.; Xie, E.; Hong, L.; Liu, X.; Liu, Z.; Lu, T.; Li, Z.; Luo, P. DDP: Diffusion model for dense visual prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 21684–21695, 2024.
- [103] Saxena, S.; Herrmann, C.; Hur, J.; Kar, A.; Norouzi, M.; Sun, D.; Fleet, D. J. The surprising effectiveness of diffusion models for optical flow and monocular depth estimation. *arXiv preprint arXiv:2306.01923*, 2023.
- [104] Yin, W.; Wang, X.; Shen, C.; Liu, Y.; Tian, Z.; Xu, S.; Sun, C.; Dou, R. DiverseDepth: Affine-invariant depth prediction using diverse data. *arXiv preprint arXiv:2002.00569*, 2020.
- [105] Lee, J. H.; Kim, C. S. Monocular depth estimation using relative depth maps. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 9721–9730, 2020.
- [106] Farooq Bhat, S.; Alhashim, I.; Wonka, P. AdaBins: Depth estimation using adaptive bins. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 4008–4017, 2021.
- [107] Bhat, S. F.; Alhashim, I.; Wonka, P. LocalBins: Improving depth estimation by learning local distributions. In: *Computer Vision – ECCV 2022. Lecture Notes in Computer Science, Vol. 13661*. Avidan, S.; Brostow, G.; Cissé, M.; Farinella, G. M.; Hassner, T. Eds. Springer Cham, 480–496, 2022.
- [108] Zhang, S.; Yang, L.; Mi, M. B.; Zheng, X.; Yao, A. Improving deep regression with ordinal entropy. *arXiv preprint arXiv:2301.08915*, 2023.
- [109] Shao, S.; Pei, Z.; Chen, W.; Chen, P. C. Y.; Li, Z. IEBins: Iterative elastic bins for monocular depth estimation and completion. *International Journal of Computer Vision* Vol. 133, No. 5, 2463–2486, 2024.
- [110] Yin, W.; Zhang, C.; Chen, H.; Cai, Z.; Yu, G.; Wang, K.; Chen, X.; Shen, C. Metric3D: Towards zero-shot metric 3D prediction from a single image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 9009–9019, 2024.
- [111] Hu, M.; Yin, W.; Zhang, C.; Cai, Z.; Long, X.; Chen, H.; Wang, K.; Yu, G.; Shen, C.; Shen, S. 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* Vol. 46, No. 12, 10579–10596, 2024.
- [112] Bochkovskii, A.; Delaunoy, A.; Germain, H.; Santos, M.; Zhou, Y.; Richter, S. R.; Koltun, V. Depth pro: Sharp monocular metric depth in less than a second. *arXiv preprint arXiv:2410.02073*, 2024.
- [113] Li, Z.; Snavely, N. MegaDepth: Learning single-view depth prediction from Internet photos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2041–2050, 2018.
- [114] Godard, C.; Mac Aodha, O.; Firman, M.; Brostow, G. Digging into self-supervised monocular depth estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 3827–3837, 2019.
- [115] Qi, X.; Liao, R.; Liu, Z.; Urtasun, R.; Jia, J. GeoNet: Geometric neural network for joint depth and surface normal estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 283–291, 2018.
- [116] Fu, X.; Yin, W.; Hu, M.; Wang, K.; Ma, Y.; Tan, P.; Shen, S.; Lin, D.; Long, X. GeoWizard: Unleashing the diffusion priors for 3D geometry estimation from a single image. In: *Computer Vision – ECCV 2024. Lecture Notes in Computer Science, Vol. 15080*. Leonardis, A.; Ricci, E.; Roth, S.; Russakovsky, O.; Sattler, T.; Varol, G. Eds. Springer Cham, 241–258, 2025.
- [117] Eftekhari, A.; Sax, A.; Malik, J.; Zamir, A. Omnidata:



- A scalable pipeline for making multi-task mid-level vision datasets from 3D scans. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 10766–10776, 2022.
- [118] Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [119] Song, Z.; Wang, Z.; Li, B.; Zhang, H.; Zhu, R.; Liu, L.; Jiang, P. T.; Zhang, T. DepthMaster: Taming diffusion models for monocular depth estimation. *arXiv preprint arXiv:2501.02576*, 2025.
- [120] Xu, G.; Ge, Y.; Liu, M.; Fan, C.; Xie, K.; Zhao, Z.; Chen, H.; Shen, C. What matters when repurposing diffusion models for general dense perception tasks? *arXiv preprint arXiv:2403.06090*, 2024.
- [121] Lavreniuk, M.; Bhat, S. F.; Müller, M.; Wonka, P. EVP: Enhanced visual perception using inverse multi-attentive feature refinement and regularized image-text alignment. In: *Computer Vision – ECCV 2024. Lecture Notes in Computer Science, Vol. 15623*. Del Bue, A.; Canton, C.; Pont-Tuset, J.; Tommasi, T. Eds. Springer Cham, 206–225, 2025.
- [122] Carreira, J.; Gokay, D.; King, M.; Zhang, C.; Rocco, I.; Mahendran, A.; Keck, T. A.; Heyward, J.; Koppula, S.; Pot, E.; et al. Scaling 4D representations. *arXiv preprint arXiv:2412.15212*, 2024.
- [123] Bhat, S. F.; Birkl, R.; Wofk, D.; Wonka, P.; Müller, M. ZoeDepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023.
- [124] Viola, M.; Qu, K.; Metzger N.; Ke, B.; Becker, A.; Schindler, K.; Obukhov, A. Marigold-DC: Zero-shot monocular depth completion with guided diffusion. *arXiv preprint arXiv:2412.13389*, 2024.
- [125] Lin, H.; Peng, S.; Chen, J.; Peng, S.; Sun, J.; Liu, M.; Bao, H.; Feng, J.; Zhou, X.; Kang, B. Prompting depth anything for 4K resolution accurate metric depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 17070–17080, 2025.
- [126] Yang, H.; Huang, D.; Yin, W.; Shen, C.; Liu, H.; He, X.; Lin, B.; Ouyang, W.; He, T. Depth any video with scalable synthetic data. *arXiv preprint arXiv:2410.10815*, 2024.
- [127] Chen, Z.; Sun, X.; Wang, L.; Yu, Y.; Huang, C. A deep visual correspondence embedding model for stereo matching costs. In: Proceedings of the IEEE International Conference on Computer Vision, 972–980, 2016.
- [128] Chang, J. R.; Chen, Y. S. Pyramid stereo matching network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 5410–5418, 2018.
- [129] Guo, X.; Yang, K.; Yang, W.; Wang, X.; Li, H. Group-wise correlation stereo network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 3268–3277, 2020.
- [130] Zhang, F.; Prisacariu, V.; Yang, R.; Torr, P. H. S. GA-net: Guided aggregation net for end-to-end stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 185–194, 2020.
- [131] Shen, Z.; Dai, Y.; Song, X.; Rao, Z.; Zhou, D.; Zhang, L. PCW-net: Pyramid combination and warping cost volume for stereo matching. In: *Computer Vision – ECCV 2022. Lecture Notes in Computer Science, Vol. 13692*. Avidan, S.; Brostow, G.; Cissé, M.; Farinella, G. M.; Hassner, T. Eds. Springer Cham, 280–297, 2022.
- [132] Badki, A.; Troccoli, A.; Kim, K.; Kautz, J.; Sen, P.; Gallo, O. Bi3D: Stereo depth estimation via binary classifications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1597–1605, 2020.
- [133] Xu, H.; Zhang, J. AANet: Adaptive aggregation network for efficient stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1956–1965, 2020.
- [134] Yang, M.; Wu, F.; Li, W. WaveletStereo: Learning wavelet coefficients of disparity map in stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 12882–12891, 2020.
- [135] He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778, 2016.
- [136] Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K. Q. Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2261–2269, 2017.
- [137] Guo, W.; Li, Z.; Yang, Y.; Wang, Z.; Taylor, R. H.; Unberath, M.; Yuille, A.; Li, Y. Context-enhanced stereo transformer. In: *Computer Vision – ECCV 2022. Lecture Notes in Computer Science, Vol. 13692*. Avidan, S.; Brostow, G.; Cissé, M.; Farinella, G. M.; Hassner, T. Eds. Springer Cham, 263–279, 2022.
- [138] Liu, Z.; Li, Y.; Okutomi, M. Global occlusion-aware transformer for robust stereo matching. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 3523–3532, 2024.



- [139] Xu, H.; Zhang, J.; Cai, J.; Rezatofighi, H.; Yu, F.; Tao, D.; Geiger, A. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 45, No. 11, 13941–13958, 2023.
- [140] Su, Q.; Ji, S. ChiTransformer: Towards reliable stereo from cues. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1929–1939, 2022.
- [141] Weinzaepfel, P.; Leroy, V.; Lucas, T.; Brégier, R.; Cabon, Y.; Arora, V.; Antsfeld, L.; Chidlovskii, B.; Csurka, G.; Revaud, J. CroCo: Self-supervised pre-training for 3D vision tasks by cross-view completion. *arXiv preprint* arXiv:2210.10716, 2022.
- [142] Lou, J.; Liu, W.; Chen, Z.; Liu, F.; Cheng, J. ELFNNet: Evidential local-global fusion for stereo matching. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 17738–17747, 2024.
- [143] Medsker, L.; Jain, L. C. *Recurrent Neural Networks: Design and Applications*. Boca Raton: CRC Press, 1999.
- [144] Xu, G.; Wang, X.; Ding, X.; Yang, X. Iterative geometry encoding volume for stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 21919–21928, 2023.
- [145] Xu, G.; Wang, X.; Zhang, Z.; Cheng, J.; Liao, C.; Yang, X. IGEV++: Iterative multi-range geometry encoding volumes for stereo matching. *arXiv preprint* arXiv:2409.00638, 2024.
- [146] Bartolomei, L.; Tosi, F.; Poggi, M.; Mattoccia, S. Stereo anywhere: Robust zero-shot deep stereo matching even where either stereo or mono fail. *arXiv preprint* arXiv:2412.04472, 2024.
- [147] Cheng, J.; Liu, L.; Xu, G.; Wang, X.; Zhang, Z.; Deng, Y.; Zang, J.; Chen, Y.; Cai, Z.; Yang, X. MonSter: Marry monodepth to stereo unleashes power. *arXiv preprint* arXiv:2501.08643, 2025.
- [148] Jiang, H.; Lou, Z.; Ding, L.; Xu, R.; Tan, M.; Jiang, W.; Huang, R. DEFOM-stereo: Depth foundation model based stereo matching. *arXiv preprint* arXiv:2501.09466, 2025.
- [149] Wang, X.; Yang, H.; Xu, G.; Cheng, J.; Lin, M.; Deng, Y.; Zang, J.; Chen, Y.; Yang, X. ZeroStereo: Zero-shot stereo matching from single images. *arXiv preprint* arXiv:2501.08654, 2025.
- [150] Guo, X.; Zhang, C.; Zhang, Y.; Wang, R.; Nie, D.; Zheng, W.; Poggi, M.; Zhao, H.; Ye, M.; Zou, Q.; et al. Stereo anything: Unifying zero-shot stereo matching with large-scale mixed data. *arXiv preprint* arXiv:2411.14053, 2024.
- [151] Watson, J.; Mac Aodha, O.; Turmukhambetov, D.; Brostow, G. J.; Firman, M. Learning stereo from single images. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science, Vol. 12346*. Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J.-M. Eds. Springer Cham, 722–740, 2020.
- [152] Bleyer, M.; Rhemann, C.; Rother, C. PatchMatch stereo - stereo matching with slanted support windows. In: Proceedings of the British Machine Vision Conference, 2011.
- [153] Schönberger, J. L.; Zheng, E.; Frahm, J. M.; Pollefeys, M. Pixelwise view selection for unstructured multi-view stereo. In: *Computer Vision – ECCV 2016. Lecture Notes in Computer Science, Vol. 9907*. Leibe, B.; Matas, J.; Sebe, N.; Welling, M. Eds. Springer Cham, 501–518, 2016.
- [154] Schönberger, J. L.; Frahm, J. M. Structure-from-motion revisited. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 4104–4113, 2016.
- [155] Triggs, B.; McLauchlan, P. F.; Hartley, R. I.; Fitzgibbon, A. W. Bundle adjustment: A modern synthesis. In: *Vision Algorithms: Theory and Practice, Lecture Notes in Computer Science, Vol. 1883*. Triggs, B.; Zisserman, A.; Szeliski, R. Eds. Springer Berlin Heidelberg, 298–372, 2000.
- [156] Hartley, R.; Zisserman, A. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003.
- [157] Montemerlo, M.; Thrun, S. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In: Proceedings of the AAAI Conference on Artificial Intelligence, 593–598, 2002.
- [158] Davison, A. J.; Reid, I. D.; Molton, N. D.; Stasse, O. MonoSLAM: Real-time single camera SLAM. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 29, No. 6, 1052–1067, 2007.
- [159] Tsai, R. Y.; Lenz, R. K. A new technique for fully autonomous and efficient 3D robotics hand/eye calibration. *IEEE Transactions on Robotics and Automation* Vol. 5, No. 3, 345–358, 1989.
- [160] Daniilidis, K. Hand-eye calibration using dual quaternions. *The International Journal of Robotics Research* Vol. 18, No. 3, 286–298, 1999.
- [161] Barnes, C.; Shechtman, E.; Finkelstein, A.; Goldman, D. B. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics* Vol. 28, No. 3, Article No. 65, 2009.
- [162] Wang, K.; Shen, S. MVDepthNet: Real-time multiview depth estimation neural network. In: Proceedings of the International Conference on 3D Vision, 248–257, 2018.



- [163] Yang, Z.; Ren, Z.; Shan, Q.; Huang, Q. MVS2D: Efficient multiview stereo via attention-driven 2D convolutions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8564–8574, 2022.
- [164] Kar, A.; Häne, C.; Malik, J. Learning a multi-view stereo machine. *arXiv preprint* arXiv:1708.05375, 2017.
- [165] Ji, M.; Gall, J.; Zheng, H.; Liu, Y.; Fang, L. SurfaceNet: An end-to-end 3D neural network for multiview stereopsis. In: Proceedings of the IEEE International Conference on Computer Vision, 2326–2334, 2017.
- [166] Cheng, S.; Xu, Z.; Zhu, S.; Li, Z.; Li, L. E.; Ramamoorthi, R.; Su, H. Deep stereo using adaptive thin volume representation with uncertainty awareness. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2521–2531, 2020.
- [167] Gu, X.; Fan, Z.; Zhu, S.; Dai, Z.; Tan, F.; Tan, P. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2492–2501, 2020.
- [168] Yang, J.; Mao, W.; Alvarez, J. M.; Liu, M. Cost volume pyramid based depth inference for multi-view stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 4876–4885, 2020.
- [169] Yi, H.; Wei, Z.; Ding, M.; Zhang, R.; Chen, Y.; Wang, G.; Tai, Y. W. Pyramid multi-view stereo net with self-adaptive view aggregation. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science, Vol. 12354*. Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J.-M. Eds. Springer Cham, 766–782, 2020.
- [170] Wang, F.; Galliani, S.; Vogel, C.; Pollefeys, M. IterMVS: Iterative probability estimation for efficient multi-view stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8596–8605, 2022.
- [171] Wang, X.; Zhu, Z.; Huang, G.; Qin, F.; Ye, Y.; He, Y.; Chi, X.; Wang, X. MVSTER: Epipolar transformer for efficient multi-view stereo. In: *Computer Vision – ECCV 2022. Lecture Notes in Computer Science, Vol. 13691*. Avidan, S.; Brostow, G.; Cissé, M.; Farinella, G. M.; Hassner, T. Eds. Springer Cham, 573–591, 2022.
- [172] Cao, C.; Ren, X.; Fu, Y. MVSFormer: Multi-view stereo by learning robust image features and temperature-based depth. *arXiv preprint* arXiv:2208.02541, 2022.
- [173] Xu, Y.; Shi, Z.; Wang, Y.; Chen, H.; Yang, C.; Peng, S.; Shen, Y.; Wetzstein, G. GRM: Large Gaussian reconstruction model for efficient 3D reconstruction and generation. *arXiv preprint* arXiv:2403.14621, 2024.
- [174] Wang, J.; Karaev, N.; Ruppert, C.; Novotny, D. VGGSfM: Visual geometry grounded deep structure from motion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 21686–21697, 2024.
- [175] Lu, J.; Huang, T.; Li, P.; Dou, Z.; Lin, C.; Cui, Z.; Dong, Z.; Yeung, S. K.; Wang, W.; Liu, Y. Align3R: Aligned monocular depth estimation for dynamic videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 22820–22830, 2025.
- [176] Wang, Q.; Zhang, Y.; Holynski, A.; Efros, A. A.; Kanazawa, A. Continuous 3D perception model with persistent state. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 10510–10522, 2025.
- [177] Wang, H.; Agapito, L. 3D reconstruction with spatial memory. *arXiv preprint* arXiv:2408.16061, 2024.
- [178] Duisterhof, B.; Züst, L.; Weinzaepfel, P.; Leroy, V.; Cabon, Y.; Revaud, J. MAST3R-SfM: A fully-integrated solution for unconstrained structure-from-motion. *arXiv preprint* arXiv:2409.19152, 2024.
- [179] Wang, Y.; Doersch, C.; Arandjelovi, R.; Carreira, J.; Zisserman, A. Input-level inductive biases for 3D reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 6166–6176, 2022.
- [180] Guo, H.; Zhu, H.; Peng, S.; Lin, H.; Yan, Y.; Xie, T.; Wang, W.; Zhou, X.; Bao, H. Multi-view reconstruction via SfM-guided monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 5272–5282, 2025.
- [181] Zhang, S.; Wang, J.; Xu, Y.; Xue, N.; Ruppert, C.; Zhou, X.; Shen, Y.; Wetzstein, G. FLARE: Feed-forward geometry, appearance and camera estimation from uncalibrated sparse views. *arXiv preprint* arXiv:2502.12138, 2025.
- [182] Yang, J.; Sax, A.; Liang, K. J.; Henaff, M.; Tang, H.; Cao, A.; Chai, J.; Meier, F.; Feiszli, M. Fast3R: Towards 3D reconstruction of 1000+ images in one forward pass. *arXiv preprint* arXiv:2501.13928, 2025.
- [183] Jang, W.; Weinzaepfel, P.; Leroy, V.; Agapito, L.; Revaud, J. Pow3R: Empowering unconstrained 3D reconstruction with camera and scene priors. *arXiv preprint* arXiv:2503.17316, 2025.
- [184] Li, Z.; Tucker, R.; Cole, F.; Wang, Q.; Jin, L.; Ye, V.;



- Kanazawa, A.; Holynski, A.; Snavely, N. MegaSaM: Accurate, fast, and robust structure and motion from casual dynamic videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 10486–10496, 2025.
- [185] Shriram, J.; Trevithick, A.; Liu, L.; Ramamoorthi, R. RealmDreamer: Text-driven 3D scene generation with inpainting and depth diffusion. In: Proceedings of the International Conference on 3D Vision, 882–892, 2025.
- [186] Zhang, J.; Herrmann, C.; Hur, J.; Jampani, V.; Darrell, T.; Cole, F.; Sun, D.; Yang, M. H. MonST3R: A simple approach for estimating geometry in the presence of motion. *arXiv preprint* arXiv:2410.03825, 2024.
- [187] Wang, P.; Tan, H.; Bi, S.; Xu, Y.; Luan, F.; Sunkavalli, K.; Wang, W.; Xu, Z.; Zhang, K. PF-LRM: Pose-free large reconstruction model for joint pose and shape prediction. *arXiv preprint* arXiv:2311.12024, 2023.
- [188] Zhang, H.; Li, Y.; Cao, Y.; Liu, Y.; Shen, C.; Yan, Y. Exploiting temporal consistency for real-time video depth estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 1725–1734, 2019.
- [189] Wang, Y.; Luo, L.; Shen, X.; Mei, X. DynOcc: Learning single-view depth from dynamic occlusion cues. In: Proceedings of the International Conference on 3D Vision, 514–523, 2020.
- [190] Chen, S.; Guo, H.; Zhu, S.; Zhang, F.; Huang, Z.; Feng, J.; Kang, B. Video depth anything: Consistent depth estimation for super-long videos. *arXiv preprint* arXiv:2501.12375, 2025.
- [191] Sayed, M.; Gibson, J.; Watson, J.; Prisacariu, V.; Firman, M.; Godard, C. SimpleRecon: 3D reconstruction without 3D convolutions. *arXiv preprint* arXiv:2208.14743, 2022.
- [192] Curless, B.; Levoy, M. A volumetric method for building complex models from range images. In: Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, 303–312, 1996.
- [193] Wang, J.; Wang, P.; Long, X.; Theobalt, C.; Komura, T.; Liu, L.; Wang, W. NeuRIS: Neural reconstruction of indoor scenes using normal priors. *arXiv preprint* arXiv:2206.13597, 2022.
- [194] Yu, Z.; Peng, S.; Niemeyer, M.; Sattler, T.; Geiger, A. MonoSDF: Exploring monocular geometric cues for neural implicit surface reconstruction. *arXiv preprint* arXiv:2206.00665, 2022.
- [195] Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; Ng, R. NeRF: Representing scenes as neural radiance fields for view synthesis. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science*, Vol. 12346. Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J.-M. Eds. Springer Cham, 405–421, 2020.
- [196] Kerbl, B.; Kopanas, G.; Leimkuehler, T.; Drettakis, G. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* Vol. 42, No. 4, Article No. 139, 2023.
- [197] Guo, J.; Deng, N.; Li, X.; Bai, Y.; Shi, B.; Wang, C.; Ding, C.; Wang, D.; Li, Y. StreetSurf: Extending multi-view implicit surface reconstruction to street views. *arXiv preprint* arXiv:2306.04988, 2023.
- [198] Miao, S.; Huang, J.; Bai, D.; Qiu, W.; Liu, B.; Geiger, A.; Liao, Y. Efficient depth-guided urban view synthesis. In: *Computer Vision – ECCV 2024. Lecture Notes in Computer Science*, Vol. 15088. Leonardis, A.; Ricci, E.; Roth, S.; Russakovsky, O.; Sattler, T.; Varol, G. Eds. Springer Cham, 90–107, 2025.
- [199] Fan, Z.; Cong, W.; Wen, K.; Wang, K.; Zhang, J.; Ding, X.; Xu, D.; Ivanovic, B.; Pavone, M.; Pavlakos, G.; et al. InstantSplat: Sparse-view Gaussian splatting in seconds. *arXiv preprint* arXiv:2403.20309, 2024.
- [200] Zhou, H.; Lin, L.; Wang, J.; Lu, Y.; Bai, D.; Liu, B.; Wang, Y.; Geiger, A.; Liao, Y. HUGSIM: A real-time, photo-realistic and closed-loop simulator for autonomous driving. *arXiv preprint* arXiv:2412.01718, 2024.
- [201] Miao, S.; Huang, J.; Bai, D.; Yan, X.; Zhou, H.; Wang, Y.; Liu, B.; Geiger, A.; Liao, Y. EVolSplat: Efficient volume-based Gaussian splatting for urban view synthesis. *arXiv preprint* arXiv:2503.20168, 2025.
- [202] Yan, Y.; Lin, H.; Zhou, C.; Wang, W.; Sun, H.; Zhan, K.; Lang, X.; Zhou, X.; Peng, S. Street Gaussians: Modeling dynamic urban scenes with Gaussian splatting. In: *Computer Vision – ECCV 2024. Lecture Notes in Computer Science*, Vol. 15131. Leonardis, A.; Ricci, E.; Roth, S.; Russakovsky, O.; Sattler, T.; Varol, G. Eds. Springer Cham, 156–173, 2025.
- [203] Sandström, E.; Tateno, K.; Oechsle, M.; Niemeyer, M.; Van Gool, L.; Oswald, M. R.; Tombari, F. SplatSLAM: Globally optimized RGB-only SLAM with 3D Gaussians. *arXiv preprint* arXiv:2405.16544, 2024.
- [204] Pan, Y.; Zhong, X.; Jin, L.; Wiesmann, L.; Popovi, M.; Behley, J.; Stachniss, C. PINGS: Gaussian splatting meets distance fields within a point-based implicit neural map. *arXiv preprint* arXiv:2502.05752, 2025.
- [205] Matsuki, H.; Murai, R.; Kelly, P. H. J.; Davison, A. J. Gaussian splatting SLAM. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 18039–18048, 2024.
- [206] Mao, Y.; Yu, X.; Zhang, Z.; Wang, K.; Wang, Y.;



- Xiong, R.; Liao, Y. NGEL-SLAM: Neural implicit representation-based global consistent low-latency SLAM system. In: Proceedings of the IEEE International Conference on Robotics and Automation, 6952–6958, 2024.
- [207] OpenAI. Video generation models as world simulators. 2024. Available at <https://openai.com/index/video-generation-models-as-world-simulators/>
- [208] Hong, W.; Ding, M.; Zheng, W.; Liu, X.; Tang, J. CogVideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022.
- [209] Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; Zhang, X.; Feng, G.; et al. CogVideoX: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.
- [210] Rigter, M.; Gupta, T.; Hilmkil, A.; Ma, C. AVID: Adapting video diffusion models to world models. *arXiv preprint arXiv:2410.12822*, 2024.
- [211] Zhang, Q.; Zhai, S.; Bautista, M. A.; Miao, K.; Toshev, A.; Susskind, J.; Gu, J. World-consistent video diffusion with explicit 3D modeling. *arXiv preprint arXiv:2412.01821*, 2024.
- [212] Kang, B.; Yue, Y.; Lu, R.; Lin, Z.; Zhao, Y.; Wang, K.; Huang, G.; Feng, J. How far is video generation from world model: A physical law perspective. *arXiv preprint arXiv:2411.02385*, 2024.
- [213] Xu, Z.; Xu, Y.; Yu, Z.; Peng, S.; Sun, J.; Bao, H.; Zhou, X. Representing long volumetric video with temporal Gaussian hierarchy. *ACM Transactions on Graphics* Vol. 43, No. 6, Article No. 171, 2024.
- [214] Xu, Z.; Peng, S.; Lin, H.; He, G.; Sun, J.; Shen, Y.; Bao, H.; Zhou, X. 4K4D: Real-time 4D view synthesis at 4K resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 20029–20040, 2024.
- [215] Xu, Z.; Xie, T.; Peng, S.; Lin, H. EasyVolcap: Accelerating neural volumetric video research. In: Proceedings of the SIGGRAPH Asia Technical Communications, Article No. 22, 2023.




Zhen Xu is a fourth-year Ph.D. student in computer science at Zhejiang University, advised by Profs. Xiaowei Zhou and Sida Peng. He obtained his bachelor degree in computer science from Zhejiang University in 2022. His current research focuses on 3D/4D neural reconstruction and rendering, and

volumetric video.



degree from Tongji University in 2020.

Hongyu Zhou is currently pursuing his Ph.D. degree at Zhejiang University, working under the guidance of Prof. Yiyi Liao at the X-Dim Lab. Previously, he was a master student at Zhejiang University supervised by Prof. Deng Cai and Prof. Xiaofei at the State Key Lab of CAD&CG. He received his bachelor



engineering from Zhejiang University in 2018. He received the 2024 CCF Outstanding Doctoral Dissertation Award and was selected as an 2022 Apple Scholar in AI/ML.

Sida Peng is an assistant professor (ZJU-100 Young Professor) at Zhejiang University. He received his Ph.D. degree from the College of Computer Science and Technology, Zhejiang University in 2023, supervised by Prof. Xiaowei Zhou and Prof. Hujun Bao, and obtained his bachelor degree in information



Snively at Cornell University.

Haotong Lin is a fourth-year Ph.D. student in computer science at Zhejiang University, advised by Profs. Xiaowei Zhou and Sida Peng. He obtained his bachelor degree in computer science from Zhejiang University in 2021. In the summer of 2022, he had the opportunity to engage in a highly enjoyable collaboration with Noah



Haoyu Guo is a research scientist at Shanghai Artificial Intelligence Laboratory. He obtained his Ph.D. degree in computer science from Zhejiang University in 2025. His research interests include 3D vision, spatial intelligence, and world models.



Jiahao Shao is a student researcher at X-D Lab advised by Yiyi Liao. He is also lucky to have collaborated with Matteo Poggi. He obtained his B.Eng. degree in automation from Zhejiang University in 2024.



Peishan Yang is a master student in computer science at Zhejiang University. His research interests focus on 3D reconstruction and scene understanding.






Qinglin Yang is currently a master student in the School of Software Engineering, Zhejiang University, advised by Dr. Sida Peng and Dr. Xiaowei Zhou. He received his B.E. degree from the School of Remote Sensing and Information Engineering, Wuhan University in, 2024. His research interests include 3D foundation models for perception, understanding, and generation.



Sheng Miao is a Ph.D. student at Zhejiang University. His research interests include 3D reconstruction and generative models in autonomous driving.



Xingyi He is a fourth-year Ph.D. student in computer science at Zhejiang University, advised by Prof. Xiaowei Zhou. He received his bachelor degree from Huazhong University of Science and Technology in 2021. His research interests include 3D computer vision, robotics, image matching, and 3D reconstruction.



Yifan Wang is an M.S. student at Zhejiang University. His research interests include 3D vision, dynamic reconstruction, and neural rendering.




Yue Wang is a professor of control science and engineering at Zhejiang University. He received his Ph.D. degree from Zhejiang University in 2016, advised by Prof. Rong Xiong. His research focuses on persistent robot autonomy using AI, machine learning, and computer vision. He has developed various robotic systems, some of which have been applied in industry.



Ruizhen Hu is a distinguished professor and deputy director of the Visual Computing Research Center at Shenzhen University. She received her Ph.D. degree in applied mathematics from Zhejiang University in 2015, and was previously an assistant researcher at SIAT. Her research focuses on computer graphics and embodied AI, especially 3D modeling and agent interaction.



Yiyi Liao is an assistant professor at Zhejiang University, leading the X-D Lab. She was previously a postdoctoral researcher in the Autonomous Vision Group at the University of Tübingen and the Max Planck Institute for Intelligent Systems. She received her Ph.D. degree from Zhejiang University in 2018 and her B.S. degree from Xi'an Jiaotong University in 2013. Her research interests include 3D computer vision, scene understanding, and 3D generative models.



Xiaowei Zhou is a professor in the College of Computer Science and the State Key Laboratory of CAD&CG at Zhejiang University. Before joining ZJU in 2017, he was a postdoctoral researcher at the GRASP Lab, University of Pennsylvania. His research focuses on 3D computer vision, including 3D reconstruction, pose estimation, motion capture, scene understanding, and their applications in mixed reality and robotics.



Hujun Bao is a professor in the State Key Laboratory of CAD&CG and the College of Computer Science and Technology, Zhejiang University. He received his B.Sc. degree (1987) and Ph.D. degree (1993) in mathematics and applied mathematics from Zhejiang University. He leads the 3D graphics computing group, focusing on geometric computing, 3D visual computing, and real-time rendering. He has contributed to systems such as VisioniX, ACTS, and 2D-to-3D video conversion.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

To submit a manuscript, please go to <https://jcvm.org>.



清华大学出版社
Tsinghua University Press

Available on
IEEE Xplore®