

SDGE: Stereo Guided Depth Estimation for 360° Camera Sets

Jialei Xu¹, Wei Yin², Dong Gong³, Junjun Jiang¹, Xianming Liu^{1*}

Abstract—Depth estimation is a critical technology in autonomous driving, and multi-camera systems are often used to achieve a 360° perception. These 360° camera sets often have limited or low-quality overlap regions, making multi-view stereo methods infeasible for the entire image. Alternatively, monocular methods may not produce consistent cross-view predictions. To address these issues, we propose the Stereo Guided Depth Estimation (SGDE) method, which enhances depth estimation of the full image by explicitly utilizing multi-view stereo results on the overlap. We suggest building virtual pinhole cameras to resolve the distortion problem of fisheye cameras and unify the processing for the two types of 360° cameras. For handling the varying noise on camera poses caused by unstable movement, the approach employs a self-calibration method to obtain highly accurate relative poses of the adjacent cameras with minor overlap. These enable the use of robust stereo methods to obtain a high-quality depth prior in the overlap region. This prior serves not only as an additional input but also as pseudo-labels that enhance the accuracy of depth estimation methods and improve cross-view prediction consistency. The effectiveness of SGDE is evaluated on one fisheye camera dataset, Synthetic Urban, and two pinhole camera datasets, DDAD and nuScenes. Our experiments demonstrate that SGDE is effective for both supervised and self-supervised depth estimation, and highlight the potential of our method for advancing autonomous driving technology. Our project page is at <https://github.com/JialeiXu/SGDE>.

I. INTRODUCTION

Depth estimation is a fundamental technology for 3D scene perception in various computer vision applications, such as robotics and autonomous driving [1], [2], [3], [4]. Benefiting from the advent of deep neural networks (DNNs), many deep learning based methods have been proposed to estimate the depth from monocular cues [5], [6], [7] or/and multi-view stereo cues [8], [9], [10]. In recent real-world applications, intelligent agents (*e.g.*, autonomous driving cars) are equipped with 360° multi-camera sets including pinhole cameras [11], [12] or fisheye cameras [13]. The 360° multi-camera sets bring a more reliable 360° field of view perception of the environment, but also bring new challenges on how to properly and fully utilize them in depth estimation.

This work was supported by National Natural Science Foundation of China under Grants 92270116 and 62071155.

¹Jialei Xu, Junjun Jiang, and Xianming Liu are with the School of Computer Science and Technology, Harbin Institute of Technology, Harbin, 150001, China, E-mail: xujialei@stu.hit.edu.cn, {jiangjunjun, csxm}@hit.edu.cn. *Corresponding author.

²Wei Yin is with the University of Adelaide, Australia. E-mail: yvanwy@outlook.com.

³Dong Gong is with the School of Computer Science and Engineering, The University of New South Wales (UNSW), Sydney, Australia. E-mail: edgong01@gmail.com.

The 360° multi-view camera sets can provide powerful geometric cues for reliable 3D modeling. Instead of only using them as training supervision [14], [15], [16], the strong stereo information should be properly used in inference/testing to directly guide the depth estimation for robustness in the real world. Many multi-view stereo depth estimation methods [17], [18] may pave the way for that. But they are non-ideal and inefficient for 360° depth estimation in realistic applications (*e.g.*, autonomous driving and robotics) that require immediate and efficient depth prediction for online response in the system. Although 360° multi-view camera sets can provide full coverage of the scene, the overlap region of adjacent cameras is small (for surrounding pinhole cameras) [12], [11], or the overlapped regions are with low-quality/sparse signals (for surrounding fisheye cameras) [13]. On the other hand, in realistic scenarios, owing to the unstable mechanical structures during cars' movement, the pre-calibrated camera extrinsic parameters (*i.e.* cameras' translation and rotation to the car's ego system) may fail during movement, as shown in Fig. 2.

Existing 360° depth estimation methods [12], [15], [19], [20] mainly focus on implicitly using the texture consistency information in the overlapping areas. Many methods [14], [15], [16] only use 360° geometric cues to provide supervision in training (*e.g.*, self-supervised depth estimation methods) and cannot guarantee the multi-view consistency in overlapping areas during inference, which cannot fully use the geometric guidance in the 360° images. On the other hand, they mainly focus on pinhole cameras [12], [11], and the studies on 360° fisheye cameras [13], [19], [20] overly rely on synthetic datasets. For both of them, existing works focus on the ideal scenarios and may ignore or overlook the characteristics of the realistic 360° images, *e.g.*, low-quality overlapping area and the pose measurement noise, as discussed above.

In this paper, we address the above issues by proposing a general pipeline named Stereo Guided Depth Estimation (SGDE) for 360° depth estimation, which *explicitly* exploits the geometric cues in the overlapping regions in both inference and training. In our approach, we propose to use classical stereo-matching methods [21], [22] to infer depth in the overlapping regions as a *depth prior* for more accurate depth estimation with DNNs, which can be very efficient and thus practical for the real-world application scenarios of the 360° depth estimation, comparing to previous MVS methods [18], [17]. By triangulating the depth prior on the overlap, we can obtain dense depth maps on the overlapping regions as prior and use them as input guidance (usable in inference) for DNNs or/and pseudo ground truth to provide

supervision (in training). The explicit stereo guidance enables more powerful guidance and robust predictions than other implicit techniques [14], [15], [16]. Specifically, to unify the process for different types of cameras, we propose to build virtual pinhole cameras for the fisheye cameras via camera model transformation. For the 360° fisheye cameras with significant overlap but serious image quality problems (i.e., uneven pixel density and distortion), the accuracy of the high-quality unsupervised depth prior exceeds multiple supervised methods. For the 360° pinhole cameras with small overlaps, our method can significantly improve the accuracy, particularly mitigating depth consistency issues on the overlap. Moreover, to handle the varying measurement noise/error on camera poses, we propose to dynamically optimize and update the camera pose in the proposed framework, which guarantees the effectiveness of stereo rectification in realistic applications (as discussed in Sec. III-B). We summarize our contributions in the following:

- We target the 360° view depth estimation problem in autonomous driving scenarios, and propose a general pipeline to solve the 360° fisheye and pinhole camera sets with explicit stereo guidance, called Stereo Guided Depth Estimation (SGDE).
- We propose to build virtual pinhole cameras to solve the distortion problem of uneven pixel density and serious distortion of fisheye cameras in the geometry method, which unifies the process of the two types of 360° camera sets.
- To handle the varying measurement noise/caused by unstable movement in realistic scenarios, we introduce the geometric loop constraint to calibrate camera poses by surrounding cameras with small overlap, which achieves higher accuracy than the ground truth and ensures the effectiveness of SGDE.
- We present experimental results on Synthetic Urban [13] and two popular real multi-camera datasets DDAD [12] and nuScenes [11] showing that SGDE can effectively improve the performance of SOTA models in both supervised and self-supervised scenarios.

II. RELATED WORK

Monocular Depth Estimation. Monocular depth estimation aims at predicting the depth map from a single image. Eigen *et al.* [23] propose the pioneered approach which inspired many follow-up works. Subsequent works propose to directly regress the continuous depth map from the aggregation of the information in an image [24], [25], [26], [27], and continuously improve the accuracy. [28], [29] formulate depth prediction into classification, and predict the depth class of each pixel to get the depth map. Due to the heavy cost of obtaining high-quality ground-truth from depth sensors, self-supervised depth prediction [30] has attracted much attention in recent years.

Multi-Camera Depth Estimation. Stereo matching has been an active research area for decades [31], [22]. Traditional methods [32], [31] focus on looking for local matching points and global optimization to exploit spatial context.

Recently, deep learning has brought big improvements to stereo matching. Lipson *et al.* [21] introduce RAFT-Stereo, a new deep architecture for rectified stereo based on the optical flow network RAFT [10]. As the demand for 360° surround perception in autonomous driving continues to rise, there is an increased focus on depth reconstruction by multi-camera in surrounding views. In Synthetic Urban dataset [13], the camera rig is virtually implemented with four fisheye cameras, and [13], [19], [20] rely on the large areas of overlap and perfect depth ground-truth to predict panorama depth. However, the high-quality ground truth is hard to obtain in real data. In surrounding pinhole camera setups, there is only a small overlap between adjacent cameras such as the DDAD dataset [12] and nuScenes dataset [11]. Guizilini *et al.* [14] first propose the self-supervised learning of scale-aware network based on mono2 in the multi-camera setting. SurroundDepth [15] employs a joint network to process all the surrounding views and proposes a cross-view transformer to effectively fuse the information from multiple views. Xu *et al.* [16] formulate the depth estimation as a weighted combination of depth basis. However, these methods only implicitly learn the information in the overlapping area, which is less effective.

Multi-Camera Calibration. There are two main approaches to the calibration problem: 1) *Chart-based calibration* assumes that an object with precisely known geometry is presented in all images, and computes the camera parameters consistent with a set of correspondences between the features defining the chart and their observed image projections. 2) *Structure from Motion (SfM)* [33], which estimates both the scene shape and the camera parameters. In SfM, the intrinsic parameters are either known in a prior or recovered posterior through auto-calibration. Typically, camera parameters and scene points are computed by non-linear optimization during bundle adjustment [34]. The selection of feature correspondences is a key step of both approaches. SIFT [35] and ORB [36] are arguably the most successful hand-crafted local features and are widely adopted in many 3D vision tasks. In the surrounding pinhole camera case, insufficient overlap area between cameras can negatively impact the performance of camera calibration in existing methods. To address this issue, we propose to use the relative pose of adjacent cameras to form a loop in geometry.

III. THE PROPOSED METHOD

Multiple cameras, such as pinhole cameras or fisheye cameras, are usually equipped in real-world applications (e.g., autonomous cars) to obtain surround 360° views. We aim to solve this 360° view depth estimation task by proposing a stereo-guided pipeline, as shown in Fig. 1. We observe that all adjacent cameras have some overlaps, which can provide confident multi-view observation of the scene at the same time. We propose to obtain the depth prior from the overlapping areas and use it explicitly to enhance the performance of depth prediction networks. This motivates us to explicitly use the stereo-matching methods to obtain the depth in overlapping regions and then use them to supervise

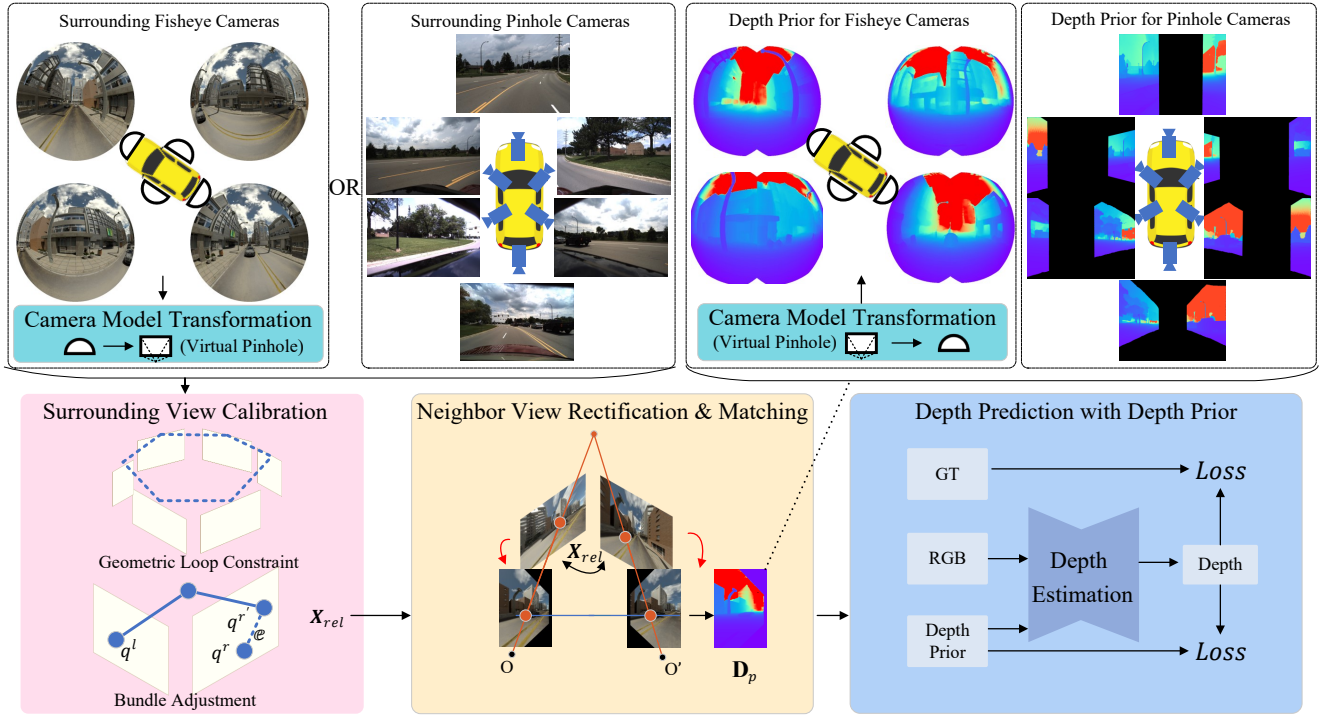


Fig. 1. **The pipeline of Stereo Guided Depth Estimation (SGDE).** For the 360° camera setup, we first propose the geometry loop constraint to optimize the surrounding camera poses, which ensures the effectiveness of the subsequent geometry scheme. Then, we use mature stereo rectification and matching algorithms to obtain the depth prior of the overlapping regions. Finally, depth prior is used as both an extra input and the supervisor signal during training of the depth estimation network.




Fig. 2. Overlaid stereo rectified images using the relative pose between adjacent two cameras in DDAD dataset [12]. The rectified images calculated according to the “ground-truth” pose provided by the dataset have a large error in the vertical direction, which indicates the inaccuracy of the dataset-provided parameters.

training or as input. However, in autonomous driving scenarios, the adjacent cameras are not normal stereo cameras, such as completely different focal lengths, noisy camera extrinsic parameters, and two fisheye cameras. We propose several modules to solve them, including camera model transformation, surrounding view calibration, and neighbor view rectification and matching.

A. Depth Prior Construction

1) *Camera Model Transformation:* We first propose to build virtual pinhole cameras to solve the problem of uneven pixel density and serious distortion of surrounding fisheye cameras in the stereo-matching methods, which can unify the process for pinhole and fisheye cameras. Specifically, the field of view in one fisheye camera is evenly divided into left and right parts corresponding to two virtual pinhole cameras.

Using one virtual camera as an example, we manually set the intrinsic parameters \mathbf{K}_{pin} and extrinsic parameters $\mathbf{X}_{pin} = \begin{pmatrix} \mathbf{R}_{pin} & \mathbf{t}_{pin} \\ 0 & 1 \end{pmatrix}$ for the virtual pinhole camera. The fisheye image is transformed into the pinhole image by applying the following warping transformation. In the fisheye image, we suppose a point located at $\mathbf{p} = (\Phi, \Theta)$ in 2D spherical coordinates, and is transformed to the spherical projection of the space point $\hat{\mathbf{p}}$:

$$\hat{\mathbf{p}} = [\sin(\Phi)\cos(\Theta), \sin(\Phi)\sin(\Theta), \cos(\Phi)]^T, \quad (1)$$

where $\hat{\mathbf{p}}$ is the unit directional vector. Then, $\hat{\mathbf{p}}$ is projected to the point $\mathbf{p}_{pin} = (u, v)$ in image coordinates of the pinhole camera as following: $\mathbf{p}_{pin} = \mathbf{K}_{pin}\mathbf{R}_{pin}^{-1}(\mathbf{R}_f\hat{\mathbf{p}} + \mathbf{t}_f - \mathbf{t}_{pin})$, where $\mathbf{X}_f = \begin{pmatrix} \mathbf{R}_f & \mathbf{t}_f \\ 0 & 1 \end{pmatrix}$ is the extrinsic parameters of fisheye cameras. We also perform linear interpolation sampling to compute the virtual pinhole image. Here, we obtain the pinhole images of the 360° views from fisheye cameras and corresponding calibrations.

2) *Neighbor View Rectification and Matching:* We can then derive the relative pose between each pair of adjacent cameras in the same frame from datasets or the module of camera model transformation (Sec. III-A.1), and use standard stereo methods to reconstruct a reliable depth prior on the overlap. The overall process of depth prior construction includes three steps: i) *Stereo rectification*, which projects overlapping images into a common image plane and decides the overlapping area. ii) *Stereo matching*, which computes the depth of the overlap region based on known camera intrinsic and extrinsic. iii) *Depth back-projection*, which back-projects the depth from the common image plane to




Fig. 3. Depth prior in Synthetic Urban [13] dataset.

the origin camera plane. We call it depth prior \mathbf{D}_p . We use an off-the-shelf stereo matching algorithm, e.g., SGBM [22] or Raft-Stereo [21], to obtain \mathbf{D}_p . Note that we do not re-train Raft-Stereo, but directly use the model trained on the Middlebury dataset [37] for depth reconstruction. SGBM is a non-parametric method.

B. Surrounding View Calibration

Though the process of depth prior construction seems straightforward, it heavily relies on the accuracy of camera calibrations. The relative poses between cameras provided in common autonomous driving datasets are very noisy or even completely fail. This is because the vibration during car driving is likely to change the location or orientation of the camera. As shown in Fig. 2, we overlay the rectified images using the relative poses provided in the dataset [12]. The same object is not on the horizontal epipolar line, which indicates the inaccuracy of the dataset-provided calibration parameters. Such noisy calibrations will cause failure in stereo matching and depth calculation. To test this point, we evaluate the depth prior computed by the calibration parameters from datasets. As shown in Tab. II, using the ground-truth camera calibrations (provided in the datasets) results in unacceptable depth reconstructions for both datasets used in the experiments. Therefore, we argue that the ground-truth camera parameters, specifically, the camera poses provided in common autonomous driving datasets are not accurate enough for the depth prior construction.

To address this problem, we need to optimize “ground-truth” camera poses between adjacent cameras. The process of high-precision stereo camera calibration is similar. By extracting correspondences in the overlap region between cameras and accumulating them over multiple frames, the relative pose can be computed by the method of epipolar geometry. However, we find this strategy to be ineffective in DDAD [12] and nuScenes [11], as the overlap region is too small. Insufficient overlap results in the failure of two-view geometry algorithms frequently as shown in Tab. II. Even with the SOTA stereo camera calibration method HOM [38], the error of the depth prior computed by the optimized pose is still significant. Hence, we propose to establish the geometric loop constraint (GLC) in 360° views to calibrate the surrounding cameras’ pose synchronously. We use the relative poses of the 360° view cameras to form a loop, which has a stronger restraint ability in the optimization process than the ordinary stereo calibration schemes.

Specifically, we optimize the poses for all cameras within a video sequence. The overall process for our calibration method contains 3 steps: 1) *Correspondence extraction*, which finds correspondences between all images within a small number of frames. 2) *Pose parametrization*: we apply a temporal consistency assumption, and assume that the relative pose between cameras remains consistent in a short period of time. 3) *Bundle adjustment*: the pose of all cameras is optimized simultaneously, making it possible to create a geometric loop constraint on the 360° views, and compute the final poses of all cameras based on the extracted correspondences and pose parametrization.

Given a video sequence with T frames and M cameras, we denote \mathbf{I}_t^m as the image from camera m in frame t . We adopt LoFTR [39] to extract candidate matching from all possible $T \times T \times M$ pairs of images from adjacent cameras, and then filter outlier matching and get the final correspondences $\{\mathbf{q}_{t_1}^{m_1}, \mathbf{q}_{t_2}^{m_2}\}$ for each image pair $\{\mathbf{I}_{t_1}^{m_1}, \mathbf{I}_{t_2}^{m_2}\}$ using Graph-Cut RANSAC [40]. Denote the forward-facing camera id $m = 1$, we first parametrize the pose \mathbf{X}_t^1 of the forward-facing camera in each frame t . For other cameras, we only optimize the relative pose \mathbf{X}_{rel}^m from each camera to the forward-facing camera. The absolute pose \mathbf{X}_t^m of camera $m > 1$ in frame t can thus be computed by:

$$\mathbf{X}_t^m = \mathbf{X}_{rel}^m \mathbf{X}_t^1. \quad (2)$$

We optimize all surrounding cameras synchronously to establish geometry loop constraint, and use standard bundle adjustment [33] to optimize the poses:

$$\underset{\{\mathbf{x}_t^1\}, \{\mathbf{X}_{rel}^m\}, \{\hat{\mathbf{q}}_i\}}{\text{minimize}} \sum_{m=1}^M \sum_{t=1}^T \sum_{i=1}^N |\mathbf{q}_{t,i}^m - \pi(\mathbf{K}^m, \mathbf{X}_t^m, \hat{\mathbf{q}}_i)|, \quad (3)$$

where \mathbf{K}^m is the intrinsic matrix of camera m , $\mathbf{q}_{t,i}^m \in \mathbf{q}_t^m$ is the 2D pixel location of an observed 3D point $\hat{\mathbf{q}}_i$ in image \mathbf{I}_t^m , N is the number of observed 3D point $\hat{\mathbf{q}}_i$, and $\pi(\cdot)$ is the projection function from a 3D point to a camera plane. Intuitively, the above equation jointly optimizes all the $T + M - 1$ poses along with the 3D location of all correspondences.

C. Depth Prediction with Depth Prior

With the virtual camera set up to solve the image quality problem of the fisheye camera, and obtaining accurate camera poses in real data, we can now produce reliable depth prior \mathbf{D}_p as described in Sec. III-A. Then we use it to enhance the depth prediction.

First of all, \mathbf{D}_p can be used as a part of the input to the model. Denote $f(\cdot|\boldsymbol{\theta})$ as the depth estimation model parametrized by $\boldsymbol{\theta}$, we predict the depth map \mathbf{D} :

$$\mathbf{D} = f(\mathbf{I}, \mathbf{D}_p|\boldsymbol{\theta}), \quad (4)$$

where \mathbf{I} is the color image. Due to its simplicity, this method can be applied to both supervised or self-supervised depth estimation models, and provide consistent improvements, as shown later in Sec. IV.

Due to the defects of the supervision signal, e.g., the sparsity of Lidar and the dynamic object problem in temporal




Fig. 4. Visualization of self-supervised depth prediction results on 6 cameras at the same frame in DDAD [12]. Top (RGB): RGB pictures of six cameras. The red box denotes the overlapping region of each camera. The second row: results of the baseline R18 [41]. The third row: results of SurroundDepth [15]. The last row (R18 (i & L1)): results of baseline R18 [41] applied with depth prior as input and L1 supervision.

TABLE I

PERFORMANCE OF SGDE IN SYNTHETIC URBAN DATASET [20].

UNSUPERVISED DENOTES THE DEPTH PRIOR IS INFERRED BY RAFT-STEREO [10] WITHOUT RE-TRAINING. SGBM [22] AND RAFT (RAFT-STEREO [21]) DENOTES THE DEPTH PRIOR IS COMPUTED BY THE CORRESPONDING METHODS.

Supervision	Method	>1 ↓	>3 ↓	>5 ↓	MAE ↓	RMS ↓
Unsupervised	RAFT-Stereo	19.29	11.48	9.57	1.22	3.48
Ground-truth	PSMNet [44]	92.67	31.45	21.32	4.33	7.76
	DispNet [45]	39.02	21.12	14.47	2.37	4.85
	Omni [19]	14.10	4.01	2.64	0.82	2.45
	+RAFT input	13.48	3.79	2.52	0.78	2.35
	+SGBM input	12.22	3.62	2.35	0.73	2.29

self-supervised signal, the loss on the depth prior can further provide extra supervision signals. Specifically, in real scenes, we add extra loss terms into the training objectives during model training via

$$\underset{\theta}{\text{minimize}} \quad L_{ori}(\mathbf{D}, \mathbf{S}_{ori}) + \lambda L_{dp}(\mathbf{D}, \mathbf{D}_p), \quad (5)$$

where $L_{ori}(\cdot)$ and \mathbf{S}_{ori} are the original loss and supervision signal of the model, and $L_{dp}(\cdot)$ is the extra supervision provided by the depth prior. We set the coefficient λ to 0.005 in all experiments. This method is inspired by [42], where the ground-truth depth is both given as inputs and supervision signals to form a depth completion task and boost the depth prediction accuracy. We experiment in Sec. IV with different realizations of $L_{dp}(\cdot)$, including the L1 loss, the virtual normal loss [7] and the ranking loss [43], which all provide effective improvements and verify our hypothesis.

IV. EXPERIMENTS

A. Experimental Setup

Datasets: We use three datasets for the experiment Synthetic Urban [13], DDAD [12], and nuScenes [11]. Synthetic Urban is virtually implemented with the camera rig by four fisheye

TABLE II

EFFECT OF POSE OPTIMIZATION ON THE ACCURACY OF THE DEPTH PRIOR. THE “GROUND TRUTH” CAMERA POSES PROVIDED IN THE DATASETS (GT) ARE NOT ACCURATE ENOUGH FOR GEOMETRY METHODS, RESULTING IN HIGH DEPTH RECONSTRUCTION ERROR. THE SOTA STEREO CALIBRATION METHOD HOM [38] CAN NOT REDUCE THE POSE ERROR EFFECTIVELY. USING GEOMETRIC LOOP CONSTRAINT (GLC) TO CALIBRATE CAMERAS CAN EFFECTIVELY REDUCE THE ERROR OF DEPTH PRIOR THAN NOT USING IT (W/O GLC).

Dataset	Pose	Abs Rel ↓	Sq Rel ↓	RMSE ↓	$\delta_{1.25}$ ↑
DDAD [12]	GT	0.885	8.392	24.184	0.439
	HOM [38]	0.599	6.381	18.209	0.498
	w/o GLC	0.474	5.827	15.230	0.565
	GLC	0.142	2.331	10.103	0.803
nuScenes [11]	GT	0.940	6.129	13.113	0.410
	HOM [38]	0.885	5.827	12.609	0.446
	w/o GLC	0.779	5.523	10.993	0.511
	GLC	0.498	3.931	9.102	0.597

cameras, and the dataset is rendered using Blender. The FOV of each camera is 220° . DDAD and nuScenes are urban driving datasets captured by six calibrated cameras time-synchronized with relatively small overlaps.

Implementation details: We use RAFT-Stereo [21] as the default stereo matching method. We use original hyper-parameters when applying SGDE to existing models. In DDAD and nuScenes, we optimize the camera pose once per sequence by the first 7 frames and the optimized pose is applied to the whole sequence.

B. Main Result

As shown in Tab. I, II, the accuracy of unsupervised depth prior surpasses that of some Lidar supervised methods in Synthetic Urban and DDAD datasets. We also show the visualization of the depth prior in Fig. 3. As shown in Tab. I, III, IV, and V, SGDE can be widely used in fisheye

TABLE III

PERFORMANCE OF SGDE FOR SUPERVISED DEPTH ESTIMATION IN NUSCENES. L1 DENOTES THE L1 LOSS ON DEPTH PRIOR. DEPTH PRIOR IS COMPUTED BY RAFT-STEREO [21].

Method	Abs Rel↓	Sq Rel↓	RMSE↓	$\delta_{1.25}$ ↑
Baseline (R18)[41]	0.113	1.266	3.332	0.914
+ input	0.095	1.124	2.937	0.936
+ input & L1	0.089	1.118	2.925	0.939
Adabins [29]	0.125	0.910	4.115	0.866
+ input	0.112	0.861	3.700	0.882
+ input & L1	0.104	0.847	3.568	0.892

TABLE IV

ABS REL OF SELF-SUPERVISED DEPTH RESULTS IN DDAD. 'w/o D_p ' DENOTES THE RESULTS OF EACH METHOD WITHOUT DEPTH PRIOR. L1 DENOTES THE L1 LOSS ON DEPTH PRIOR. SGBM [22] AND RAFT (RAFT-STEREO [21]) DENOTES THE DEPTH PRIOR IS COMPUTED BY THE CORRESPONDING METHODS.

Method	W/O D_p	+ input	+ L1	+input & L1
SC-Depth [46]	0.210	-	-	-
SC-Depth + Ours(RAFT)	-	0.185	0.201	0.178
PackNet [12]	0.215	-	-	-
PackNet + Ours(RAFT)	-	0.188	0.206	0.177
FSM [14]	0.202	-	-	-
FSM + Ours(RAFT)	-	0.185	0.191	0.178
MCDP [16]	0.193	-	-	-
MCDP + Ours(RAFT)	-	0.180	0.187	0.177
SD [15]	0.200	-	-	-
SD + Ours(RAFT)	-	0.181	0.190	0.175
Baseline (R18) + SGBM	0.220	0.191	0.213	0.184
Baseline (R18) + RAFT	0.220	0.186	0.209	0.178

and pinhole cameras. And in *both* self-supervised and supervised tasks, SGDE significantly and consistently improves the baselines and SOTA models. Applying the depth prior as an extra input (+ input) provides effective improvements. Adding depth prior as extra supervision signals (VN/RL/L1) can further improve the performance. We also show in Fig. 4 and 5 the qualitative results of different methods in DDAD dataset [12]. With the help of SGDE, the Resnet-18 base model produces sharper and more accurate results than its baseline and SOTA method.

C. Ablation Studies

Pose Optimization Matters: We evaluate the accuracy of the depth prior calculated from the different calibration methods. As shown in Tab. II, using inaccurate “ground-truth” poses (GT) resulted in poor depth prior. Because of the small overlap area between cameras, the existing SOTA stereo camera calibration method HOM [38] can not effectively reduce the error. With geometric loop constraint (GLC), our surround view calibration method can effectively reduce the error of the depth prior. In order to ablate the role of the proposed geometry loop constraint in our calibration method, we perform our calibration scheme on two adjacent cameras,




Fig. 5. **Multi-view point clouds from predicted depth.** SGDE (ours) produces much more consistent point clouds across views. View 3 shows the result of the reconstruction of the ground, our result of different views is on the same plane, while SOTA method (SurroundDepth [15]) is not.

TABLE V

PERFORMANCE OF SELF-SUPERVISED DEPTH ESTIMATION RESULTS IN NUSCENES. VN, RL, L1 DENOTE THE VIRTUAL NORMAL LOSS [7], RANKING LOSS [43] AND L1 LOSS ON DEPTH PRIOR.

Method	Abs Rel↓	Sq Rel↓	RMSE↓	$\delta_{1.25}$ ↑
Baseline (R18) [41]	0.308	3.460	8.014	0.659
+ input	0.289	3.212	7.249	0.691
+ input & VN	0.270	3.102	7.058	0.704
+ input & RL	0.276	3.029	7.010	0.700
+ input & L1	0.268	3.190	6.949	0.705

which cannot build the geometry loop constraint. As shown in Tab. II, experimental results show that compared with our method using geometry loop constrain (GLC), the error of depth prior is significantly increased due to the lack of that (w/o GLC). This result is consistent in both datasets and proves that the geometry loop constraint is indeed effective in the surrounding view calibration method.

Stereo Matching Methods: We evaluate the effect of the different stereo matching algorithms in SGDE. We compare the hand-crafted method SGBM [22] with learning based method RAFT-Stereo [21]. As shown in Tab. I and IV, SGBM also works in our SGDE pipeline, and this non-parametric approach imposes little burden on practical applications, but the benefits are great.

V. CONCLUSION

In this paper, we propose Stereo Guided Depth Estimation for 360°Camera Sets (SGDE), a general pipeline for depth prediction from multiple overlapping cameras. We demonstrate that explicitly rectifying and stereo matching adjacent images give us accurate stereo depth priors. The core of the contribution is an optimization algorithm to obtain the more accurately calibrated camera pose than “ground-truth”. Adding the stereo depth prior to all the existing depth estimation model can consistently improve the accuracy and the cross-view geometric consistency. Finally, we quantitatively and qualitatively demonstrate the effectiveness of our method on three public datasets.

TABLE VI
THE ABS REL ERROR OF STEREO-MATCHING (RAFT) UNDER
DIFFERENT MATCHING SCHEMES.

Method	Static Scenes	Dynamic Objects
Spatio-temporal Match	0.145	0.438
Match in the same frame (Ours)	0.144	0.136

1 Benchmark and Evaluation Metrics

We evaluate the generality of SGDE by applying it to both pinhole and fisheye camera models in both self-supervised and supervised depth estimation tasks. In Synthetic Urban dataset, we follow [19] to evaluate the performance of SGDE in existing supervised SOTA models [19], [44], [45]. In DDAD and nuScenes datasets, following [14] in the self-supervised task, we choose Mono2 [41] with Resnet-18 (R18) as the baseline, and include two monocular methods PackNet [12] and SC-Depth [46]. We also evaluate three multi-camera depth estimation methods FSM [14], MCDP [16], and SurroundDepth [15].

2 More Ablation Studies

1) *Depth Consistency in Overlap Regions*: The consistency of the predicted depth on the overlap regions is an important metric for downstream tasks such as 3D reconstruction. We show in Fig. 6 the point cloud predicted by the Resnet-18 based self-supervised method on six cameras in DDAD and purposely zoom in on a car in the overlapping area. In Fig. 6b and 6c, the car is heavily distorted, which is caused by the cross-view depth prediction inconsistency. After applying SGDE, the car becomes smooth and consistent in Fig. 6d.

2) *Challenges in Establishing Spatio-Temporal Matches*: It is a common practice to establish spatio-temporal matches between different cameras at different times to increase the overlap area. For example, the back-right camera at time T matches with the front-right camera at time $T - K$, for some value of K . However, this strategy faces significant limitations in autonomous driving applications. Spatio-temporal matches assumes that the entire scene is static, but dynamic objects (*i.e.*, pedestrians, bicycles and vehicles) have different positions at different frames, which causes the method of MVS to fail. And dynamic objects are important to observe in autonomous driving applications. To test this idea, we use the image of one camera at the T frame, and the image of the adjacent camera at the $T + 10$ frame to perform stereo matching in DDAD dataset. As shown in Tab. VI, spatio-temporal matching introduces significant errors in the depth estimation of dynamic objects (pedestrians, bicycles, and vehicles), but our method is not affected by them.

REFERENCES

[1] K. Cheng, X. Long, W. Yin, J. Wang, Z. Wu, Y. Ma, K. Wang, X. Chen, and X. Chen, "Uc-nerf: Neural radiance field for under-

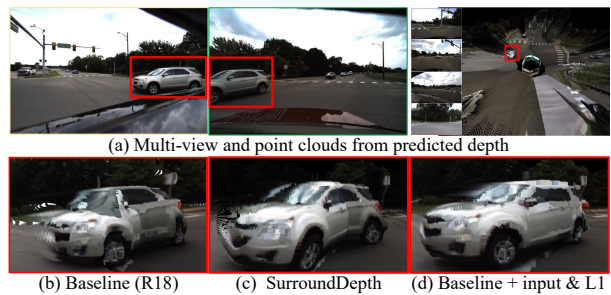


Fig. 6. Self-supervised predicted point cloud result comparison on DDAD dataset. (a) are multiple camera views in the same frame and the predicted surrounding point clouds. The car appearing in the overlapping area is enlarged in (b), (c) and (d) by different methods.

calibrated multi-view cameras in autonomous driving," *arXiv preprint arXiv:2311.16945*, 2023.

- [2] W. Yin, J. Zhang, O. Wang, S. Niklaus, L. Mai, S. Chen, and C. Shen, "Learning to recover 3d scene shape from a single image," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 204–213.
- [3] W. Yin, C. Zhang, H. Chen, Z. Cai, G. Yu, K. Wang, X. Chen, and C. Shen, "Metric3d: Towards zero-shot metric 3d prediction from a single image," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 9043–9053.
- [4] K. Cheng, X. Long, K. Yang, Y. Yao, W. Yin, Y. Ma, W. Wang, and X. Chen, "Gaussianpro: 3d gaussian splatting with progressive propagation," in *Forty-first International Conference on Machine Learning*, 2024.
- [5] J. Cheng, W. Yin, K. Wang, X. Chen, S. Wang, and X. Yang, "Adaptive fusion of single-view and multi-view depth for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 10 138–10 147.
- [6] R. Li, T. Fischer, M. Segu, M. Pollefeys, L. Van Gool, and F. Tombari, "Know your neighbors: Improving single-view reconstruction via spatial vision-language reasoning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 9848–9858.
- [7] W. Yin, Y. Liu, C. Shen, and Y. Yan, "Enforcing geometric constraints of virtual normal for depth prediction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [8] R. Li, D. Gong, W. Yin, H. Chen, Y. Zhu, K. Wang, X. Chen, J. Sun, and Y. Zhang, "Learning to fuse monocular and multi-view cues for multi-frame depth estimation in dynamic scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 539–21 548.
- [9] J. Wu, R. Li, H. Xu, W. Zhao, Y. Zhu, J. Sun, and Y. Zhang, "Gomvs: Geometrically consistent cost aggregation for multi-view stereo," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 20 207–20 216.
- [10] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *European conference on computer vision*. Springer, 2020, pp. 402–419.
- [11] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscnets: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [12] V. Guizilini, R. Ambrus, S. Pillai, A. Raventos, and A. Gaidon, "3d packing for self-supervised monocular depth estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2485–2494.
- [13] C. Won, J. Ryu, and J. Lim, "Sweepnet: Wide-baseline omnidirectional depth estimation," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6073–6079.
- [14] V. Guizilini, I. Vasiljevic, R. Ambrus, G. Shakhnarovich, and A. Gaidon, "Full surround monodepth from multiple cameras," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5397–5404, 2022.
- [15] Y. Wei, L. Zhao, W. Zheng, Z. Zhu, Y. Rao, G. Huang, J. Lu, and J. Zhou, "Surrounddepth: Entangling surrounding views

- for self-supervised multi-camera depth estimation,” *arXiv preprint arXiv:2204.03636*, 2022.
- [16] J. Xu, X. Liu, Y. Bai, J. Jiang, K. Wang, X. Chen, and X. Ji, “Multi-camera collaborative depth prediction via consistent structure estimation,” *arXiv preprint arXiv:2210.02009*, 2022.
- [17] J. Yang, W. Mao, J. M. Alvarez, and M. Liu, “Cost volume pyramid based depth inference for multi-view stereo,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4877–4886.
- [18] X. Gu, Z. Fan, S. Zhu, Z. Dai, F. Tan, and P. Tan, “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*, 2020, pp. 2495–2504.
- [19] C. Won, J. Ryu, and J. Lim, “Omnimvs: End-to-end learning for omnidirectional stereo matching,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 8987–8996.
- [20] —, “End-to-end learning for omnidirectional stereo matching with uncertainty prior,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 11, pp. 3850–3862, 2020.
- [21] L. Lipson, Z. Teed, and J. Deng, “Raft-stereo: Multilevel recurrent field transforms for stereo matching,” in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 218–227.
- [22] H. Hirschmüller, “Accurate and efficient stereo processing by semiglobal matching and mutual information,” in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, vol. 2. IEEE, 2005, pp. 807–814.
- [23] D. Eigen, C. Puhrsch, and R. Fergus, “Depth map prediction from a single image using a multi-scale deep network,” *Advances in neural information processing systems*, vol. 27, 2014.
- [24] F. Liu, C. Shen, G. Lin, and I. Reid, “Learning depth from single monocular images using deep convolutional neural fields,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 10, pp. 2024–2039, 2015.
- [25] J. H. Lee, M.-K. Han, D. W. Ko, and I. H. Suh, “From big to small: Multi-scale local planar guidance for monocular depth estimation,” *arXiv preprint arXiv:1907.10326*, 2019.
- [26] M. Song, S. Lim, and W. Kim, “Monocular depth estimation using laplacian pyramid-based depth residuals,” *IEEE transactions on circuits and systems for video technology*, vol. 31, no. 11, pp. 4381–4393, 2021.
- [27] I. Alhashim and P. Wonka, “High quality monocular depth estimation via transfer learning,” *arXiv preprint arXiv:1812.11941*, 2018.
- [28] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao, “Deep ordinal regression network for monocular depth estimation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2002–2011.
- [29] S. F. Bhat, I. Alhashim, and P. Wonka, “Adabins: Depth estimation using adaptive bins,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4009–4018.
- [30] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, “Unsupervised learning of depth and ego-motion from video,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1851–1858.
- [31] D. Scharstein and R. Szeliski, “A taxonomy and evaluation of dense two-frame stereo correspondence algorithms,” *International journal of computer vision*, vol. 47, no. 1, pp. 7–42, 2002.
- [32] H. Hirschmüller, “Stereo processing by semiglobal matching and mutual information,” *IEEE Transactions on pattern analysis and machine intelligence*, vol. 30, no. 2, pp. 328–341, 2007.
- [33] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [34] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, “Bundle adjustment—a modern synthesis,” in *International workshop on vision algorithms*. Springer, 1999, pp. 298–372.
- [35] P. C. Ng and S. Henikoff, “Sift: Predicting amino acid changes that affect protein function,” *Nucleic acids research*, vol. 31, no. 13, pp. 3812–3814, 2003.
- [36] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “Orb: An efficient alternative to sift or surf,” in *2011 International conference on computer vision*. Ieee, 2011, pp. 2564–2571.
- [37] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling, “High-resolution stereo datasets with subpixel-accurate ground truth,” in *German conference on pattern recognition*. Springer, 2014, pp. 31–42.
- [38] Y. Ling and S. Shen, “High-precision online markerless stereo extrinsic calibration,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 1771–1778.
- [39] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, “Loftr: Detector-free local feature matching with transformers,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 8922–8931.
- [40] D. Barath and J. Matas, “Graph-cut ransac,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6733–6741.
- [41] C. Godard, O. Mac Aodha, M. Firman, and G. J. Brostow, “Digging into self-supervised monocular depth estimation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 3828–3838.
- [42] M. Hu, S. Wang, B. Li, S. Ning, L. Fan, and X. Gong, “Penet: Towards precise and efficient image guided depth completion,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 656–13 662.
- [43] K. Xian, J. Zhang, O. Wang, L. Mai, Z. Lin, and Z. Cao, “Structure-guided ranking loss for single image depth prediction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [44] J.-R. Chang and Y.-S. Chen, “Pyramid stereo matching network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 5410–5418.
- [45] E. Ilg, T. Saikia, M. Keuper, and T. Brox, “Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [46] J. Bian, Z. Li, N. Wang, H. Zhan, C. Shen, M.-M. Cheng, and I. Reid, “Unsupervised scale-consistent depth and ego-motion learning from monocular video,” *Advances in neural information processing systems*, vol. 32, 2019.