

OmniStitch: Depth-Aware Stitching Framework for Omnidirectional Vision with Multiple Cameras

Omitted for Double-blind Review

ABSTRACT

Omnidirectional vision systems provide a 360-degree panoramic view, enabling full environmental awareness in various fields, such as Advanced Driver Assistance Systems (ADAS) and Virtual Reality (VR). Existing omnidirectional stitching methods rely on a single specialized 360-degree camera. However, due to hardware limitations such as high mounting heights and blind spots, adapting these methods into vehicles of varying sizes and geometries is challenging. These challenges include limited generalizability due to the reliance on predefined stitching regions for fixed camera arrays, performance degradation from distance parallax leading to large depth differences, and the absence of suitable datasets with ground truth for multi-camera omnidirectional systems. To overcome these challenges, we propose a novel omnidirectional stitching framework and publicly available dataset tailored for varying distance scenarios with multiple cameras. The framework, referred to as , consists of a Stitching Region Maximisation (SRM) module for automatic adaptation to different vehicles with multiple cameras and a Depth-Aware Stitching (DAS) module to handle depth differences caused by distance parallax between cameras. In addition, we create and release an omnidirectional stitching dataset, called which provides ground truth images that maintain the perspective of the 360-degree FOV, specifically designed for vehicle-agnostic systems. Extensive evaluations of this dataset demonstrate that our framework outperforms state-of-the-art stitching models, especially in handling varying distance parallax. The proposed dataset and code are publicly available in URL.

CCS CONCEPTS

• **Computing methodologies** → **Image processing**; **Virtual reality**; **Reconstruction**.

KEYWORDS

Omnidirectional vision, Image stitching, Omnidirectional view dataset

ACM Reference Format:

Omitted for Double-blind Review. 2018. OmniStitch: Depth-Aware Stitching Framework for Omnidirectional Vision with Multiple Cameras. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Unpublished working draft. Not for distribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted by ACM, provided that the copies are not made for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

2024-04-13 11:53. Page 1 of 1–10.

1 INTRODUCTION

Omnidirectional vision systems are becoming increasingly important in a wide range of applications, from advanced driver-assistant systems (ADAS), Virtual Reality (VR) to Unmanned Robotics and Mars exploration [8, 43, 44, 46], providing essential 360-degree environmental awareness. Traditional image stitching methods for omnidirectional vision [17, 30, 32, 51, 55] that rely on specialized 360° cameras (e.g., Samsung Gear 360¹ and Insta360²) mainly focus on single camera settings with multiple wide-angle lenses and synthesize the images from difference views overcoming large angular parallax. However, the single-camera setting is not suitable for general vehicle applications such as autonomous driving [14, 16, 47, 48, 58] due to high installation requirements [21, 26] and blind spots [15, 45], which limit their effectiveness across different vehicle sizes and designs. To provide all-around 360° coverage for vehicle-agnostic omnidirectional vision system (VA-OVS), image stitching techniques with multiple camera settings should be adopted as illustrated in Figure 1.

Addressing the technical challenges of omnidirectional image stitching in VA-OVS scenarios presents several technical difficulties. First, the existing image stitching methods are limited to inputs from 360° cameras [17, 24, 35, 51] or specific camera arrays designed for predefined [23, 56], manually crafted stitching regions. When adapting to different camera settings, manually aligning the images, transforming the viewpoint, and selecting a well-matched stitching region require redundant human efforts. Second, these methods struggle with distance parallax between different camera inputs. Our observation indicates that increased distance parallax leads to diminished stitching performance, as illustrated in Figure 1. This degradation is mainly due to the extreme depth differences across images, which prior methods overlook, resulting in errors such as ghosting and misalignment artifacts. Third, there are no data sets with ground truths for omnidirectional image stitching with dynamic distance parallax that can be applied to VA-OVS scenarios.

To address these challenges, we propose a novel omnidirectional stitching framework, called OmniStitch for dynamic distance consideration with multiple cameras. OmniStitch has two main modules: Stitching Region Maximization (SRM) and Depth-Aware Stitching (DAS). The SRM module solves the stitching region decision difficulty by simplifying the image alignment and viewpoint transformation process by statically matching half of the image for each stitching input pair. By maximizing the candidate size of image stitching, the framework can learn dynamic distance parallax between cameras and eliminate the need for pre-processing and post-processing. After the SRM process, the DAS module transforms image pairs by considering large depth differences between image pixels and extracted context features to generate a seamless

¹<https://www.samsung.com/sec/gear/360/>

²<https://www.insta360.com/>

59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116

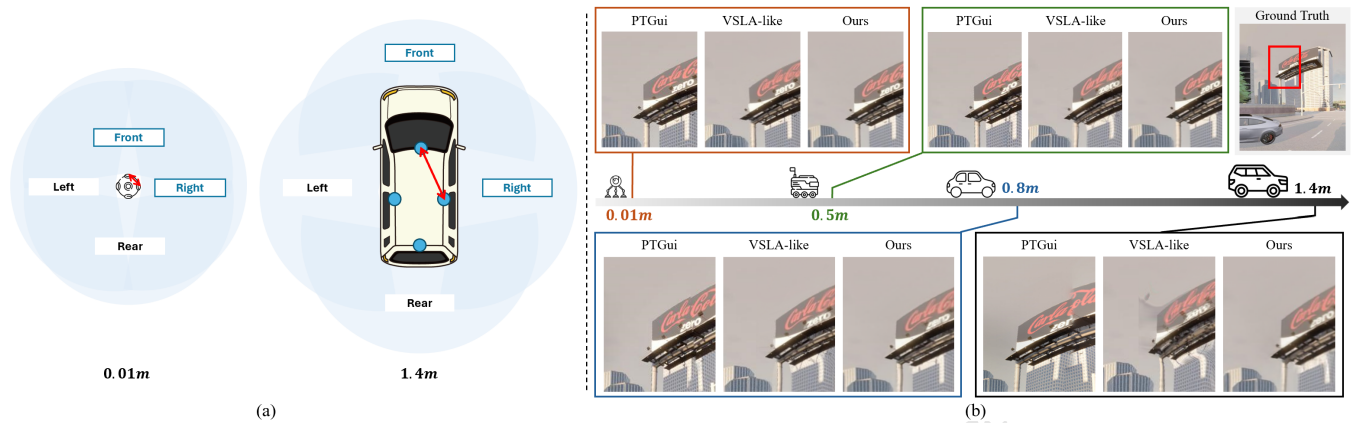


Figure 1: Examples of Vehicle Agnostic Omnidirectional Vision System (VA-OVS) (a): Inputs with dynamic distance for four-way camera configuration. (b): Comparison of the outputs of the image stitching methods with the ground truth image.

360-degree image. It enhances traditional flow estimation by incorporating additional depth information computed from 3D feature matching, and warps the images with learnable weights at the pixel level to consider large depth differences between the images.

For the challenge of dataset scarcity in the field of omnidirectional vision, we release the General-purpose Vehicle’s 360-degree (GV360) dataset, leveraging the CARLA [10] simulator to generate ground-truth images in a virtual environment. This dataset is designed with data captured from four-way multiple cameras, incorporating a wide range of dynamic distance parallax inputs and preserving a 360° FOV perspective. The use of virtual environments overcomes the limitations associated with real-world data collection, broadening the scope for future research in vehicle-agnostic omnidirectional vision systems (VA-OVS). Experimental results using the GV360 demonstrate that our framework outperforms the state-of-the-art stitching model by an average of 18.5% of PSNR and 18.0% of LPIPS for all tested distance parallax scenarios. Furthermore, our framework demonstrates superior performance on unseen datasets in qualitative results.

The main contributions are summarized as follows:

- We propose a novel end-to-end omnidirectional image stitching framework for dynamic distance parallax in multiple camera settings. It enhances generalization of stitching framework to multi-camera Vehicle-Agnostic Omnidirectional Vision System (VA-OVS) scenarios by effectively eliminating the need for a hand-crafted stitching region decision process.
- Our framework outperforms the state-of-the-art stitching methods, regardless of variations in distance parallax between cameras. The SRM module simplifies the manual process of identifying stitching regions for multiple cameras, and the DAS module effectively handles the artifacts caused by distance parallax between cameras.
- To tackle the dataset sparsity problem in omnidirectional stitching fields, we release a dataset for omnidirectional vision stitching, referred to as GV360. This dataset includes ground-truth images that capture 360-degree perspectives and simulate varying distance parallax scenarios across different vehicle types.

2 RELATED WORK

2.1 Unidirectional Stitching

Most image stitching models are designed for unidirectional vision applications, which generate a panoramic image from the viewpoint of a reference image. Unidirectional stitching methods merge two images by warping one towards the other to seamlessly stitch boundary regions [4, 6, 11, 13, 18, 25, 27, 28, 31, 36, 40, 54, 59], [60]. Various techniques were proposed to divide the image into sub-regions and estimate homography for each region, using ground and distance planes [13], grid units [28, 59, 60], triangles [27, 36], and superpixels [25]. Other enhancements have included maintaining image structure through local or global similarity transformations [4, 6, 11, 18, 31] and leveraging deep learning-based stitching models [40] for improved geometric integrity. However, these methods still suffer from projective distortion and misalignment artifacts in the case of wide parallax, due to the homography transformation process, which distorts the target image to fit the viewpoint of a reference image. Although VSLA [23] proposes to use bi-directional flow estimation to synthesize an intermediate view for wide parallax between cameras, the need for specific camera arrays and extensive pre-processing reveals the limited adaptability to various vehicles in VA-OVS scenarios.

2.2 Omnidirectional Stitching

Omnidirectional stitching combines images from all directions to create a 360-degree view, providing a perspective view of spherical maps such as the equirectangular projection (ERP) format. It is commonly used with 360-degree cameras [17, 51] or circular rig cameras [2, 35], which can produce 360-degree images by synthesizing each viewpoint. Traditional stitching methods for omnidirectional vision rely on mathematical approaches, focusing on geometric calibration through feature points [7, 17, 32, 55], calibration boxes [57], and depth map integration for seamless multi-camera setting [2], as well as advanced optical flow methods [35]. On the other hand, deep-learning approaches can significantly enhance the performance of omnidirectional stitching, such as attention mechanisms [12, 30] for matching SR to human visual perceptions and weakly-supervised

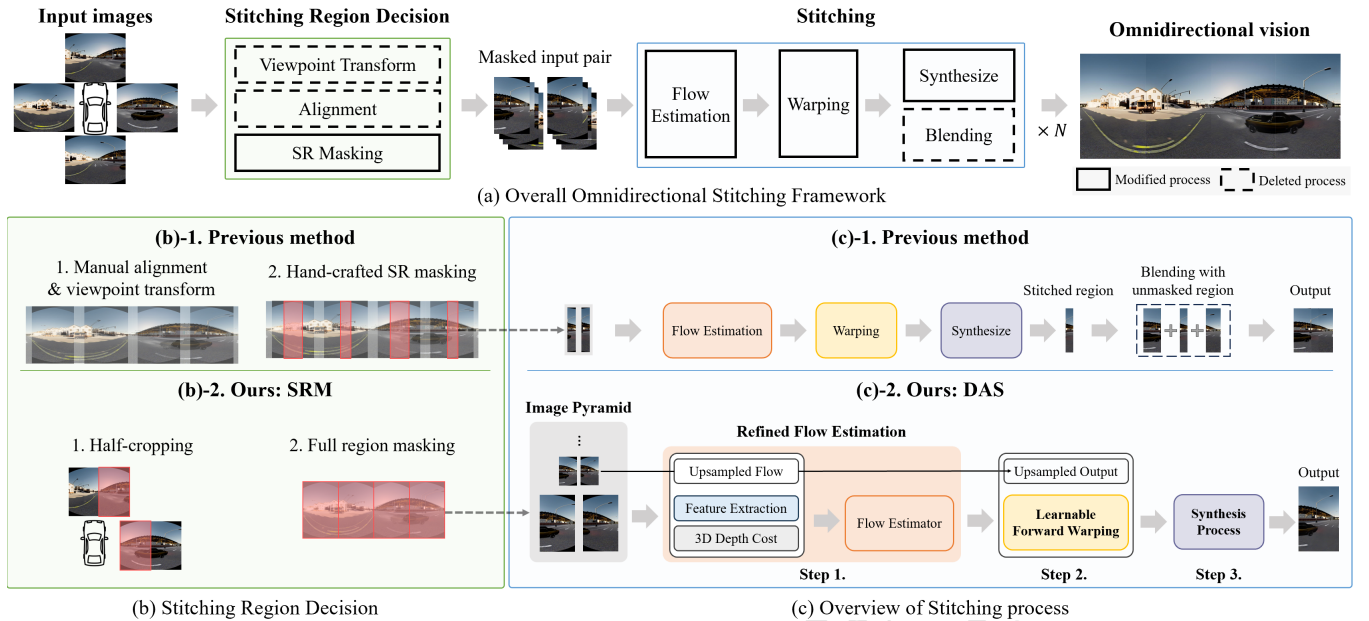


Figure 2: Overview of image stitching framework. In (a) and (c), the borderline boxes represent the modified process, and the solid lines represent the deleted process from the previous stitching framework thanks to the stitching region maximization (SRM) module. (b) and (c) are divided into the previous method and our proposed method. In (c), the orange, yellow, and purple boxes mean each of the consequent steps.

learning [51] for real-world ground truth image training. However, despite their efforts for wide parallax, they only support fixed camera formations (e.g., dual 195° FOV lenses or quad 190° FOV lenses) with a single 360° camera, and do not consider dynamic distance parallax in scenarios requiring flexible camera positioning.

2.3 Image Stitching Datasets

Datasets are crucial for the development of image stitching methods, particularly those that require ground truth (GT) or deal with parallax. View-free [37], Stitched MS-COCO [41], and DIR-D [39] provide GT for unidirectional stitching but lack parallax consideration. In contrast, datasets for large parallax such as VSLA [23], UDIS-D [38] and PDIS [22] suffer from restricted access or lack of GT, limiting their applicability. For omnidirectional stitching, CROSS [29] and WSSN [51] provide accessible GT but are confined to single camera systems, restricting the versatility of camera configurations. The GV360 dataset fills these gaps by supporting multiple camera settings for VA-OVS using virtual environments to ensure adaptability across different scenarios, thus broadening the scope for stitching algorithm development.

3 DEPTH-AWARE OMNIDIRECTIONAL STITCHING FRAMEWORK

3.1 Overview

Our goal is to integrate images from multiple cameras into a single omnidirectional image, accommodating variations in distance parallax that arise due to the diverse sizes and geometries of vehicles. In this work, we identify two requirements for VA-OVS scenarios:

1) the implementation of a multi-camera setting and 2) adaptability to dynamic distance between cameras depending on diverse vehicle types. This approach assumes a configuration with four surrounding cameras mounted on the vehicle's body. This is the minimum number of cameras in a typical scenario with multiple cameras [15, 47, 48, 58], arranged at 90° angles between them, providing an omnidirectional configuration ideal for vehicles of various shapes and sizes. We consider a dynamic distance parallax scenario where the distance between two cameras can vary from 0.01m to 1.4m, which is suitable for various vehicles.

The overall process of the omnidirectional stitching framework is illustrated in Figure 2. Given the input $\{IMG_i\}_{i=1}^N$, where N is the number of image inputs (i.e., $N = 4$), the general framework first identifies the optimal stitching regions for each pair of adjacent images. These regions are denoted as $\{IMG_{(i,r)}, IMG_{(i+1,l)}\}_{i=1}^N$, where r and l indicate the identified right and left parts of the images, respectively. These paired images undergo bidirectional flow estimation and pixel-level warping during the stitching phase, resulting in the output $\{O_i\}_{i=1}^N$, where each output

$$O_i = \text{Stitching}(IMG_{(i,r)}, IMG_{(i+1,l)})$$

is a segment of the panoramic image created by the stitching process. We simplify the notation of stitching region pairs from $\{IMG_{(i,r)}, IMG_{(i+1,l)}\}$ to $\{I_L, I_R\}$.

To ensure automatic adaptation of stitching region pairs $\{I_L, I_R\}$ to input images with dynamic distance parallax caused by multiple camera settings, we suggest using the Stitching Region Maximization (SRM) module instead of the manual process before stitching. Furthermore, the Depth-Aware Stitching (DAS) module is proposed

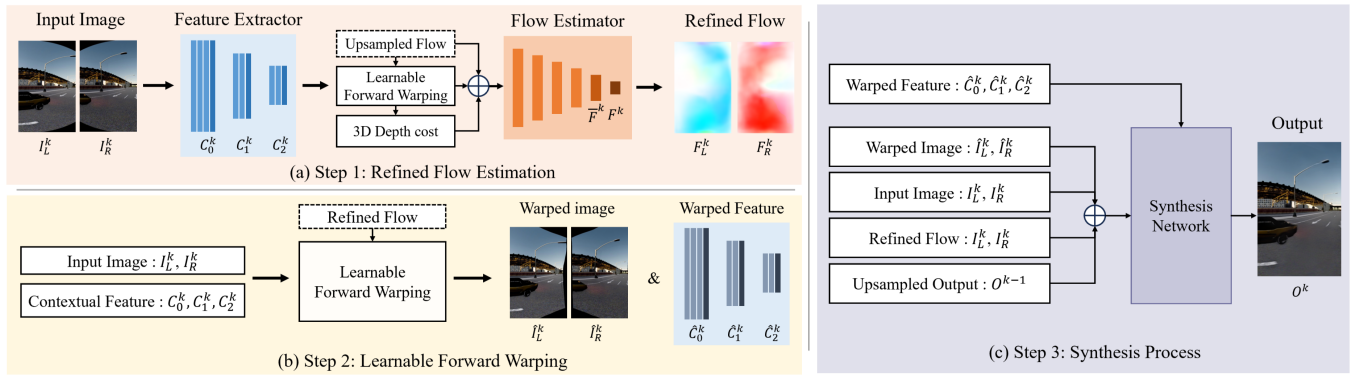


Figure 3: The step-wise details of the DAS module. The refined flow estimation process (Step 1) from input images to refined flow is illustrated in (a), while (b) and (c) present the input and output configurations for the learnable forward warping and synthesis process, respectively.

to address artifacts of synthesized outputs O_i caused by depth differences by utilizing additional depth information and an advanced warping network.

3.2 Stitching Region Maximization Module

In previous approaches, the SR decision for each input pair typically involves manually transforming the viewpoint and aligning the images or using automatic algorithms such as PTGui³. In this case, the region near the seamline of the aligned images is defined as the SR pairs. Only this region will be masked and stitched, followed by a post-processing step such as blending with other unmasked regions. Conversely, we implement half-cropping and full-region matching for each pair I_L, I_R , thereby maximizing the stitching region. This technique simplifies the SR decision process by eliminating the need for manual pre-processing (e.g., image alignment, viewpoint transformation, and stitching region decision) and post-processing (e.g., blending with unmasked regions), facilitating adaptation to various camera formations with dynamic distance.

Although the SRM allows for adaptation to various vehicles with multiple camera settings, its performance may decrease when adjusting to situations with large distance parallax. This is due to the increased candidate size for flow estimation, which can confuse the pixel-wise warping process. To handle the increased number of candidates resulting from SRM, we use a cascaded refinement strategy. It is based on the up-sampled features, flow, and synthesis results, which are fed into the next level of an image pyramid structure inspired by the frame interpolation network [19, 49, 61]. This structure organizes the stitching process into multiple levels, with each level k of the total K levels refining the input by incorporating up-sampled output from the previous level. This iterative refinement process ensures a more accurate and robust output.

3.3 Depth-Aware Stitching Module

While addressing the challenge of stitching region decisions to handle variations in distance parallax, it was observed that stitching performance decreases in scenarios with large distance parallax, as shown in Figure 1. Taking inspiration from the task of Video

³<https://ptgui.com>

Frame Interpolation (VFI) [9] tasks, which manage large motions and depth differences – challenges similar to those in large distance parallax stitching – we develop the Depth-Aware Stitching (DAS) module. The DAS module incorporates a methodological enhancement inspired by video frame interpolation models [19, 42, 52], which are tolerant of large motion. It modifies each step of the traditional image stitching process to address such challenges related to dynamic distance parallax. The module consists of three main steps: Refined Flow Estimation (Step 1), Learnable Forward Warping (Step 2), and Synthesis Network (Step 3). Each step integrates elements of the UPR-Net [19] structure adapted to the specific requirements of VA-OVS scenarios, with further details provided in Figure 3.

3.3.1 Step 1: Refined Flow Estimation. Traditionally, pixel-level flow features are extracted directly from the input pair's stitching region $\{I_L, I_R\}$ using a basic flow estimator $\mathcal{F}(\cdot)$ composed of stacked convolution layers. At each pyramid level k , traditional methods utilize outputs from the fifth and sixth layers of a stacked convolution neural network (CNN) to define flow features \bar{F}^k and flow F^k respectively as:

$$\bar{F}^k = CNN_5(I^k), F^k = CNN_6(I^k)$$

with bi-directional flow estimation.

To enhance flow estimation in the context of large distance parallax, our framework inspired by frame interpolation methods [19, 20], leveraging multi-scale feature extraction, up-sampled flow from previous pyramid level, and 3D depth costs. This approach mitigates flow estimation errors caused by large candidate sizes from stitching region maximization using up-sampled flow and brightness discrepancies caused by large parallax. Incorporating high-level contextual information via multi-scale features effectively reduces these impacts, providing additional spherical depth information and considering both explicit and implicit structural information.

Specifically, the feature extractor consists of three multi-scale stages, with each stage comprising 3, 4, and 4 convolution layers, respectively, which process both left and right images to produce context features $\{C_{(L,0)}^k, C_{(R,0)}^k\}$, $\{C_{(L,1)}^k, C_{(R,1)}^k\}$, and $\{C_{(L,2)}^k, C_{(R,2)}^k\}$ at increasing channel depths of 24, 48, and 96. These features are

Table 1: Comparison of datasets for wide parallax stitching. It includes camera settings, input types, camera formats, distance parallax, availability, and output field-of-view. The camera format contains the degree of utilized lenses.

Model	Camera setting	Input type	Camera format	Distance parallax	Availability	Output FOV
CROSS[29]	Single 360° camera	Dual 180°	Fisheye (195°)	0.01 m	Y	360°
WSSN[51]	Single 360° camera	Triple 120°	Fisheye (195°)	0.01 m	Y	360°
Quad-fisheye[7]	Single 360° camera	Quadratic 90°	Fisheye (190°)	0.01 m	N	360°
VSLA[23]	Multicamera	Triple 45°	Rectilinear (120°)	1.28 m	N	210°
GV360 (Ours)	Multicamera	Quadratic 90°	Fisheye (185°)	0.01 m - 1.4 m	Y	360°

then used to compute the 3D depth cost, accurately match features between aligned images, and adjust for spherical depth differences. The refined flow encoder integrates these depth costs with up-sampled flow, formulated as:

$$F^k = \mathcal{F}(F^{k-1} \oplus \bar{F}^{k-1} \oplus \text{Cost}(C_2^k) \oplus \bar{C}_2^k)$$

with 3D depth cost $\text{Cost}(C_2^k)$ and warped feature \bar{C}_2^k related to the current feature level.

3.3.2 Step 2: Learnable Forward Warping. Given the input image pair I_L^k, I_R^k and optical flow F_L^k, F_R^k , the traditional flow-based warping methods typically use mathematical backward warping to produce the warped image pair:

$$\{\bar{I}_L^k, \bar{I}_R^k\} = \mathcal{W}(I_L^k, I_R^k, F_L^k, F_R^k)$$

However, these methods result in issues such as object occlusion or duplication in scenarios with large depth differences. To address these challenges effectively, we employ a learnable forward warping technique inspired by the softmax splatting method [42], which enhances the handling of depth differences and reduces artifacts in the warped images. Specifically, given the input image pair $\{I_L^k, I_R^k\}$ and refined flow $\{F_L^k, F_R^k\}$, the warped left image is defined by the equation:

$$\bar{I}_L^k = \vec{w}(I_L^k, F_L^{k-1}/2, Z)$$

where \vec{w} represents the forward warping function. It utilizes half of the up-sampled flow F_L^{k-1} to find intermediate pixel positions between the left and right images, while a depth metric Z aids in accurate pixel mapping. Adopting the softmax splatting approach, the depth metric Z is used to assign appropriate weights to pixels within I_L^k for their projection onto corresponding target pixels in I_R^k , expressed as:

$$Z = \lambda \cdot v(C_L^k, (\|C_L^k - \overleftarrow{w}(C_R^k, F_R/2)\|_1))$$

where the weighting procedure is derived from backward warping \overleftarrow{w} and a scaling map λ , with v representing a U-Net for training.

The core of our learnable forward warping technique lies in its capacity to adaptively determine the source of pixels, especially in scenarios where multiple pixels may target the exact location of the image. This process is made differentiable, allowing for effective learning and optimization. This warping process is also used for feature warping \bar{C}^k during the flow estimation process as outlined in Step 1.

3.3.3 Step 3: Synthesis Process. Finally, recursively warped images \bar{I}_L^k and \bar{I}_R^k are combined using an encoder-decoder structure based on the U-Net framework [19, 20, 61] to generate an intermediate view between two images. Given synthesis network inputs, which include the current stitching sources I^k, \bar{I}^k, F^k , and the recursive output O^{k-1} from the prior pyramid level, the stitched image for each pyramid level k is formulated as:

$$O_i^k = \frac{M_L^k \odot \bar{I}_L^k + M_R^k \odot \bar{I}_R^k}{M_L^k + M_R^k} + \Delta^k$$

where M_L^k and M_R^k are the masking maps generated by the synthesis network, located in the fourth and fifth channels, respectively, and Δ^k represents a residual map contained within the first three channels. The final omnidirectional image O is obtained by smoothly concatenating the stitched segments $\{O_i\}_{i=1}^N$.

3.4 Training Loss

OmniStitch framework aims to optimize the parameters θ across the feature encoder, flow encoder \mathcal{F} , warping network \mathcal{W} , and stitching network \mathcal{S} , jointly minimizing the three distinct losses during the training phase. These include the Charbonnier [5] and VGG loss [50] for pixel-level accuracy and assessing perceptual quality, respectively. In addition, the Census loss \mathcal{L}_{cen} , known for its effectiveness in handling object occlusion [34], is used to consider large depth differences between two images. The overall training loss is expressed as:

$$\mathcal{L}_{train} = \mathcal{L}_{Char} + \alpha \mathcal{L}_{vgg} + \beta \mathcal{L}_{cen}$$

where α and β are calibrating each loss term's scale.

4 GV360 DATASETS

Our goal is to facilitate supervised learning for omnidirectional image stitching to produce precise 360°Equirectangular Projection (ERP) images, addressing the need for accurate ground-truth (GT) data and adaptability to multi-camera configurations in typical vehicles that closely mirror real-world conditions. As shown in Table 1, although several datasets exist for wide parallax image stitching [7, 23], there are still limitations in terms of public availability and provision of accurate GT images required for training supervised image stitching models. Certain omnidirectional stitching datasets [29, 51] provide GT images captured with specialized cameras, but these do not align with the requirements of the VA-OVS scenarios. To address these issues, we introduce a new method for generating datasets that can be used in multi-camera settings and

Table 2: Quantitative result of the GV360 dataset. The bold represents the best performance and the underline represents the secondary performance.

Approach	Method	GV360 dataset								
		(a) 0.01 m			(b) 0.8 m			(c) 1.4 m		
		PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
Unidirectional	APAP	11.586	0.087	0.310	11.362	0.073	0.350	10.688	0.051	0.411
	UDIS++	13.435	0.129	0.215	12.724	0.104	0.266	10.903	0.066	0.375
	LPC	13.531	0.111	0.193	13.416	0.094	0.216	12.831	0.072	0.233
	VSLA-like	<u>23.634</u>	<u>0.815</u>	<u>0.143</u>	<u>23.398</u>	<u>0.780</u>	<u>0.206</u>	<u>23.857</u>	<u>0.779</u>	<u>0.224</u>
Omnidirectional	Samsung Gear 360	14.772	0.079	0.241	14.430	0.073	0.256	14.341	0.071	0.262
	PTGui	18.885	0.227	0.229	14.832	0.153	0.284	12.597	0.061	0.327
	OmniStitch	29.146	0.911	0.108	27.147	0.871	0.169	27.710	0.870	0.192

are not specific to a particular vehicle. We present the General-Purpose Vehicle’s 360-degree (GV360) dataset generated by this method. This GV360 dataset is created using the CARLA simulator, an open-source platform that can simulate various traffic and weather conditions and camera configurations.

We render images using fisheye lenses with a wide 185-degree field of view. To satisfy vehicle-agnostic quadratic multi-camera settings [15, 47, 48, 58], we attach a total of four cameras to the upper side (U), downside (D), left (L), and right (R) of each vehicle as illustrated in Figure 1. The distance between cameras varied from 0.01m to 1.4m, adjusting to the various shapes and sizes of the vehicles. The four ground-truth images are aligned for different input pairings: front-right (RU), right-back (RD), back-left (LD), and left-front (LU). These GT images are derived by collecting the center pixels from camera inputs and continuously placing them along an ellipse. This approach minimizes distortion and closely replicates a real-world panoramic view for each side. We place 29 cameras on each of the four intermediate sides to ensure comprehensive coverage. We then un-warp the camera inputs, crop the pixels closest to the center, and align and blend the cropped pixels to reduce camera distortion, creating what we consider to be the ground truth image. The camera arrangement in this setup is inspired by the traditional rigged camera systems (*e.g.*, Google Jump [1]), where cameras are placed at 3° intervals from the center to produce a high-fidelity intermediate image. The dataset consists of 12,724 training samples and 1,284 test samples. The proposed dataset and codes are made publicly available via a provided URL.

5 EXPERIMENTS

5.1 Implementation Details

The weights in the training loss (3.4) are set to be $\alpha = 1$ and $\beta = 0.01$. The training process spans 300,000 iterations with a batch size of 10, utilizing the AdamW optimizer [33] with a weight decay of 10^{-4} . The learning rate is gradually reduced through cosine annealing from 10^{-4} to 10^{-5} . The implementation is based on Pytorch 2.1.0 with CUDA 12.1 on Ubuntu 18.04, and the training is performed on two GPUs (NVIDIA RTX 4090 and RTX A6000). To augment

the training data, we randomly crop 480×480 patches from the original data and apply random rotations, vertical flips, and channel inversion. For the image pyramid architecture, we decided to use a pyramid level of 4 and skip the flow estimation in the last level of the pyramid, instead using the cascaded flow empirically.

5.2 Baselines

To validate the performance of models in omnidirectional image stitching for general vehicle scenarios, we examine two categories of state-of-the-art models. First, we focus on unidirectional stitching methods, which predominantly handle pair-wise or linear arrays of image inputs.

- **APAP** [59] is a widely adopted stitching method, recognized for an As-Projective-As-Possible approach. It blends non-projective deviations into a globally projective framework.
- **LPC** [18] emerges as a novel solution aimed to overcome the wide parallax issue. It leverages co-planar region matching alongside global colinear structures, ensuring the preservation of geometrical integrity on both local and global scales.
- **UDIS++** [40] is a parallax-tolerant, unsupervised deep-learning model for image stitching that handles large parallax challenges and improves generalization across different datasets and resolutions.
- **VSLA-like** [23] is a representative stitching model that utilizes partial stitching regions. Since the original code is not open source, we built our own version based on the descriptions in the paper. We trained it with GV360 dataset and set the stitching regions to the same proportions as in the original model (40% of the input image).

Second, we explore omnidirectional stitching methods designed to produce 360° panoramic images.

- **Samsung Gear 360** [17] is a commercial stitching solution using dual fisheye lenses. Due to the limited reproducibility of many stitching methods, it is a benchmark for generating omnidirectional images.
- **PTGui** is a widely used commercial software for automatic stitching that employs feature points for image alignment,

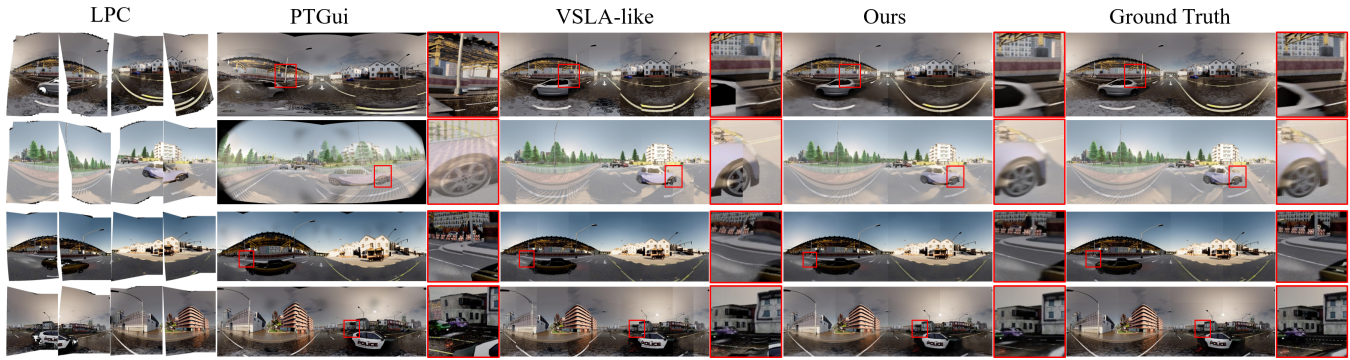


Figure 4: Qualitative results on GV360 dataset. Distance parallax = 1.4 m (1 and 2 rows), and 0.8 m (3 and 4 rows)

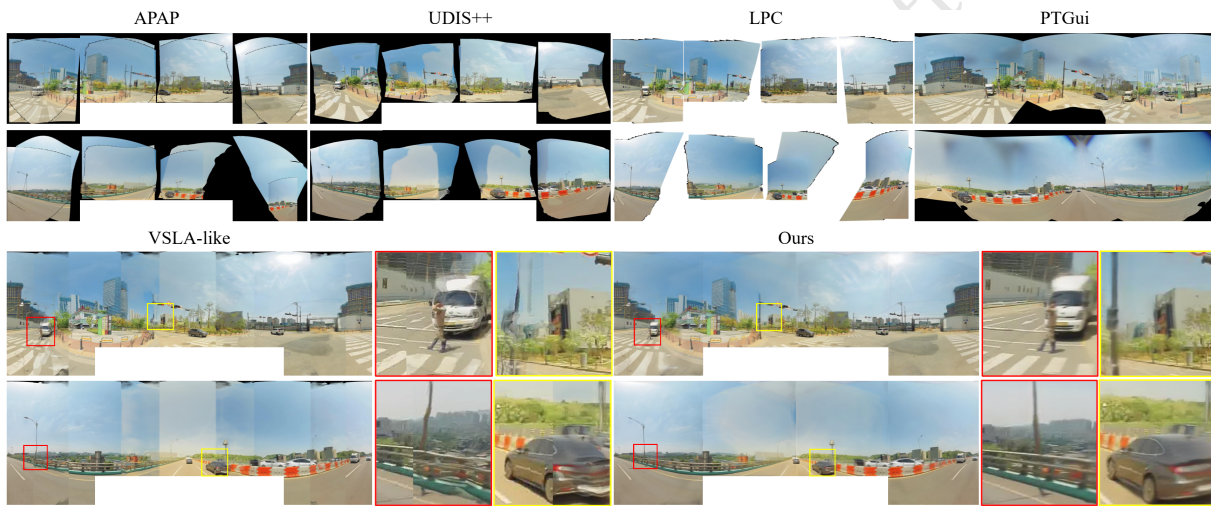


Figure 5: Qualitative results on the real dataset in large distance parallax (1.4 m between cameras) scenario.

similar to the Autostitch [3] method. It supports various camera inputs, but its automatic stitching often failed in the GV360 test set, particularly in instances with large distance parallax.

5.3 Quantitative Results

The effectiveness of our framework is evaluated by comparing it with popular image stitching methods, as outlined in 5.2. This comparison uses standard metrics such as Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM) [53], and Learned Perceptual Image Patch Similarity (LPIPS) [62]. The comprehensive results are presented in Table 2. Our experiments span various distance parallax scenarios – 0.01 m, 0.8 m, and 1.4 m – to evaluate the adaptability of the stitching methods for general vehicle applications. Our framework, equipped with the SRM module, shows superior performance over the VSLA-like method, achieving an average 18.5% improvement of PSNR, 11.7% improvement of SSIM, and 18.0% decrease of LPIPS. It verifies the benefits of applying full region adaptation with depth-aware flow refinement and a learnable warping network in stitching multiple images, thereby enhancing model adaptability for general vehicles and improving

performance in large parallax situations. Notably, our framework shows a 14.3% decrease in LPIPS over the VSLA method in the large distance parallax scenario. This significant performance gap highlights the importance of considering depth differences and reducing flow estimation errors through additional depth information and a fine-grained warping network configuration.

5.4 Qualitative results

For qualitative evaluation, we compare image stitching methods on GV360 and additional unseen datasets, focusing on the presence of visual artifacts (*e.g.*, misalignment, ghosting, and duplication) and the ability to generate natural and seamless images from a human-centric perspective for omnidirectional vision. This analysis includes three different distance parallax scenarios, similar to our quantitative multi-camera VA-OVS evaluation.

5.4.1 GV360 dataset. The qualitative results from our dataset are illustrated in Figure 4. The results reveal that LPC and UDIS++ methods are incompatible with omnidirectional stitching due to their symmetric properties, leading to significantly inferior results.

Table 3: Self-comparison with SRM and DAS module experiments on the full test dataset. ‘P’ means partial (40%) and ‘M’ means maximum stitching region.

Module	SR	IP	Cost	LFW	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
SRM	P				23.598	0.792	<u>0.215</u>
	M				<u>24.891</u>	<u>0.826</u>	0.236
DAS	M	✓			26.812	0.863	0.171
	M	✓	✓		26.999	0.867	0.164
	M	✓		✓	27.211	0.866	0.185
	M	✓	✓	✓	27.866	0.883	0.150

Commercial stitching methods such as Samsung Gear 360 and PT-Gui, although designed for 360° image stitching, struggle with large distance parallax in multi-camera settings, often resulting in duplicate and misalignment artifacts. Among these methods, VSLA-like is manually adjusted with hand-crafted image alignment and partial SR decisions to achieve more natural and seamless images. However, it is still prone to ghosting artifacts, highlighting its inability to handle depth differences. In contrast, our method effectively mitigates unintended visual artifacts and produces images that are both natural in appearance and seamlessly integrated.

5.4.2 Unseen dataset. We extend our comparison to include scenarios involving unseen data, which comprises images from real-world multi-camera settings, as shown in Figure 5. Despite the training of our framework on virtual samples from simulations, it has demonstrated robustness when applied to unseen datasets from real-world contexts. This result highlights our framework’s generalizability to a various omnidirectional vision scenarios.

5.5 Ablation Study

5.5.1 Effectiveness of Components. In the ablation study presented in Table 3, we investigate the effectiveness of different components within our proposed modules in improving the stitching quality. First, we compare the performance of the Stitching Region Maximization (SRM) module, which aims to maximize the stitching region against a traditional method focusing only on partial stitching regions. The result shows a significant improvement in performance metrics such as PSNR and SSIM, indicating numerical and structural differences in image quality. Despite these gains, extending flow estimation to the full region may cause subtle errors, potentially leading to image-wide blurring. It can result in perceptually uncomfortable outcomes, as quantified by the LPIPS metric, highlighting a trade-off between maximizing stitching area and maintaining image sharpness.

The DAS module addresses the depth difference caused by distance parallax through three key components: the image pyramid (IP), 3D depth cost (Cost), and learnable forward warping (LFW). Notably, the absence of 3D depth costs leads to a deterioration in LPIPS scores, underscoring the pivotal role of 3D depth information in the pursuit of more natural images in omnidirectional vision applications. In addition, omitting learnable forward warping reduces PSNR and SSIM. It demonstrates the effectiveness of an integral

feature of learnable warping in detecting the importance of specific object pixels in situations with substantial depth differences between images. This enables the accurate placement of objects in the foreground, significantly enhancing the image quality.

Table 4: Ablation study of training loss functions.

Training loss	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
Ours	27.866	0.883	0.150
w/o \mathcal{L}_{census}	26.774	0.859	0.176

5.5.2 Effectiveness of additional training loss. We examine the effectiveness of incorporating an additional loss term, specifically the Census loss, to improve the quality of image stitching. The comparison detailed in Table 4 demonstrates the clear advantages of integrating \mathcal{L}_{census} in our approach, as evidenced by the improvements across all the evaluated metrics. The PSNR and SSIM values indicate higher similarities to the ground truth images in pixel intensity and structural information. At the same time, the decrease in the LPIPS score shows a closer perceptual resemblance to the ground truth. This experiment demonstrates that including additional census loss effectively refines the output of our image stitching framework.

6 CONCLUSION

This paper introduces a novel omnidirectional stitching framework, referred to as OmniStitch, which addresses the challenges of dynamic distance parallax in multi-camera Vehicle-Agnostic Omnidirectional Vision Systems (VA-OVS). The framework incorporates Stitching Region Maximization (SRM) and Depth-Aware Stitching (DAS) modules to enhance stitching quality and robustness across various vehicle configurations by automating stitching region decisions and refining flow estimation with depth information to minimize artifacts. To overcome the lack of suitable datasets for VA-OVS, we provide the GV360 dataset with ground truth images that preserve 360° perspective using virtual environments. Our framework is superior in handling distance parallax variations and provides more accurate and reliable omnidirectional images. Although it currently operates with a limited number of cameras and is not completely resistant to flow estimation errors such as blurring, it provides valuable opportunities for future research to enhance the scalability and efficiency of omnidirectional vision systems for broader applications in ADAS, VR, and other fields.

REFERENCES

- [1] Robert Anderson, David Gallup, Jonathan T Barron, Janne Kontkanen, Noah Snavely, Carlos Hernández, Sameer Agarwal, and Steven M Seitz. 2016. Jump: virtual reality video. *ACM transactions on graphics (TOG)* 35, 6 (2016), 1–13.
- [2] Josep Bosch, Klemen Istenič, Nuno Gracias, Rafael Garcia, and Pere Rídao. 2019. Omnidirectional multicamera video stitching using depth maps. *IEEE Journal of Oceanic Engineering* 45, 4 (2019), 1337–1352.
- [3] Matthew Brown and David G Lowe. 2007. Automatic panoramic image stitching using invariant features. *International journal of computer vision* 74 (2007), 59–73.
- [4] Che-Han Chang, Yoichi Sato, and Yung-Yu Chuang. 2014. Shape-preserving half-projective warps for image stitching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3254–3261.
- [5] Pierre Charbonnier, Laure Blanc-Feraud, Gilles Aubert, and Michel Barlaud. 1994. Two deterministic half-quadratic regularization algorithms for computed imaging.

- In *Proceedings of 1st international conference on image processing*, Vol. 2. IEEE, 168–172.
- [6] Yu-Sheng Chen and Yung-Yu Chuang. 2016. Natural image stitching with the global similarity prior. In *European conference on computer vision*. Springer, 186–201.
- [7] Haojie Cheng, Chunxiao Xu, Jiajun Wang, and Lingxiao Zhao. 2022. Quad-fisheye Image Stitching for Monoscopic Panorama Reconstruction. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 94–109.
- [8] Ana De Abreu, Cagri Ozcinar, and Aljosa Smolic. 2017. Look around you: Saliency maps for omnidirectional images in VR applications. In *2017 Ninth International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 1–6.
- [9] Jiong Dong, Kaoru Ota, and Mianxiang Dong. 2023. Video frame interpolation: A comprehensive survey. *ACM Transactions on Multimedia Computing, Communications and Applications* 19, 2s (2023), 1–31.
- [10] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An open urban driving simulator. In *Conference on robot learning*. PMLR, 1–16.
- [11] Peng Du, Jifeng Ning, Jiguang Cui, Shaoli Huang, Xinchao Wang, and Jiaxin Wang. 2022. Geometric structure preserving warp for natural image stitching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3688–3696.
- [12] Huiyu Duan, Xionghuo Min, Wei Sun, Yucheng Zhu, Xiao-Ping Zhang, and Guangtao Zhai. 2023. Attentive deep image quality assessment for omnidirectional stitching. *IEEE Journal of Selected Topics in Signal Processing* (2023).
- [13] Junhong Gao, Seon Joo Kim, and Michael S Brown. 2011. Constructing image panoramas using dual-homography warping. In *CVPR 2011*. IEEE, 49–56.
- [14] Vitor Guizilini, Igor Vasiljevic, Rares Ambrus, Greg Shakhnarovich, and Adrien Gaidon. 2022. Full surround monodepth from multiple cameras. *IEEE Robotics and Automation Letters* 7, 2 (2022), 5397–5404.
- [15] Markus Heimberger, Jonathan Horgan, Ciarán Hughes, John McDonald, and Senthil Yogamani. 2017. Computer vision in automated parking systems: Design, implementation and challenges. *Image and Vision Computing* 68 (2017), 88–101.
- [16] Lionel Heng, Benjamin Choi, Zhaopeng Cui, Marcel Geppert, Sixing Hu, Benson Kuan, Peidong Liu, Rang Nguyen, Ye Chuan Yeo, Andreas Geiger, et al. 2019. Project autovision: Localization and 3d scene perception for an autonomous vehicle with a multi-camera system. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 4695–4702.
- [17] Tuan Ho and Madhukar Budagavi. 2017. Dual-fisheye lens stitching for 360-degree imaging. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2172–2176.
- [18] Qi Jia, ZhengJun Li, Xin Fan, Haotian Zhao, Shiyu Teng, Xinchun Ye, and Longjin Jan Latecki. 2021. Leveraging line-point consistency to preserve structures for wide parallax image stitching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12186–12195.
- [19] Xin Jin, Longhai Wu, Jie Chen, Youxin Chen, Jayoon Koo, and Cheul-hee Hahm. 2023. A unified pyramid recurrent network for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1578–1587.
- [20] Xin Jin, Longhai Wu, Guotao Shen, Youxin Chen, Jie Chen, Jayoon Koo, and Cheul-hee Hahm. 2023. Enhanced bi-directional motion estimation for video frame interpolation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 5049–5057.
- [21] Christian Kinzig, Irene Cortés, Carlos Fernández, and Martin Lauer. 2022. Real-time seamless image stitching in autonomous driving. In *2022 25th International Conference on Information Fusion (FUSION)*. IEEE, 1–8.
- [22] Hyeokjun Kweon, Hyeonseong Kim, Yoonsu Kang, Youngho Yoon, Woosong Jeong, and Kuk-Jin Yoon. 2023. Pixel-wise warping for deep image stitching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 1196–1204.
- [23] Wei-Sheng Lai, Orazio Gallo, Jinwei Gu, Deqing Sun, Ming-Hsuan Yang, and Jan Kautz. 2019. Video stitching for linear camera arrays. *arXiv preprint arXiv:1907.13622* (2019).
- [24] Hyunchul Lee, Sungmin Lee, and Okkyung Choi. 2020. Improved method on image stitching based on optical flow algorithm. *International Journal of Engineering Business Management* 12 (2020), 1847979020980928.
- [25] Kyu-Yul Lee and Jae-Young Sim. 2020. Warping residual based image stitching for large parallax. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8198–8206.
- [26] Si-Ho Lee, Bong-Ju Kim, and Seon-Bong Lee. 2021. Study on image correction and optimization of mounting positions of dual cameras for vehicle test. *Energies* 14, 16 (2021), 4857.
- [27] Jing Li, Baosong Deng, Rongfu Tang, Zhengming Wang, and Ye Yan. 2019. Local-adaptive image alignment based on triangular facet approximation. *IEEE Transactions on Image Processing* 29 (2019), 2356–2369.
- [28] Jing Li, Zhengming Wang, Shiming Lai, Yongping Zhai, and Maojun Zhang. 2017. Parallax-tolerant image stitching based on robust elastic warping. *IEEE Transactions on multimedia* 20, 7 (2017), 1672–1687.
- [29] Jia Li, Kaiwen Yu, Yifan Zhao, Yu Zhang, and Long Xu. 2019. Cross-reference stitching quality assessment for 360 omnidirectional images. In *Proceedings of the 27th ACM international conference on multimedia*. 2360–2368.
- [30] Jia Li, Yifan Zhao, Weihua Ye, Kaiwen Yu, and Shiming Ge. 2019. Attentive Deep Stitching and Quality Assessment for 360 Omnidirectional Images. *IEEE journal of selected topics in signal processing* 14, 1 (2019), 209–221.
- [31] Chung-Ching Lin, Sharathchandra U Pankanti, Karthikeyan Natesan Ramamurthy, and Aleksandr Y Aravkin. 2015. Adaptive as-natural-as-possible image stitching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1155–1163.
- [32] I-Chan Lo, Kuang-Tsu Shih, and Homer H Chen. 2021. Efficient and accurate stitching for 360° dual-fisheye images and videos. *IEEE Transactions on Image Processing* 31 (2021), 251–262.
- [33] Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017).
- [34] Simon Meister, Junhwa Hur, and Stefan Roth. 2018. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [35] Mingyuan Meng and Shaojun Liu. 2020. High-quality panorama stitching based on asymmetric bidirectional optical flow. In *2020 5th International Conference on Computational Intelligence and Applications (ICCIA)*. IEEE, 118–122.
- [36] A Moussa and Naser El-Sheimy. 2016. A fast approach for stitching of aerial images. *The international archives of the photogrammetry, remote sensing and spatial information sciences* 41 (2016), 769–774.
- [37] Lang Nie, Chunyu Lin, Kang Liao, Meiqin Liu, and Yao Zhao. 2020. A view-free image stitching network based on global homography. *Journal of Visual Communication and Image Representation* 73 (2020), 102950.
- [38] Lang Nie, Chunyu Lin, Kang Liao, Shuaicheng Liu, and Yao Zhao. 2021. Unsupervised deep image stitching: Reconstructing stitched features to images. *IEEE Transactions on Image Processing* 30 (2021), 6184–6197.
- [39] Lang Nie, Chunyu Lin, Kang Liao, Shuaicheng Liu, and Yao Zhao. 2022. Deep rectangling for image stitching: a learning baseline. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5740–5748.
- [40] Lang Nie, Chunyu Lin, Kang Liao, Shuaicheng Liu, and Yao Zhao. 2023. Parallax-tolerant unsupervised deep image stitching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7399–7408.
- [41] Lang Nie, Chunyu Lin, Kang Liao, and Yao Zhao. 2020. Learning edge-preserved image stitching from large-baseline deep homography. *arXiv preprint arXiv:2012.06194* (2020).
- [42] Simon Niklaus and Feng Liu. 2020. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5437–5446.
- [43] Janice Pan, Vikram Appia, Jesse Villarreal, Lucas Weaver, and Do-Kyoung Kwon. 2017. Rear-stitched view panorama: A low-power embedded implementation for smart rear-view mirrors on vehicles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 20–29.
- [44] Mark W Powell, Jeffrey S Norris, MA Vona, Paul G Backes, and Justin V Wick. 2005. Scientific visualization for the Mars Exploration Rovers. In *Proceedings of the 2005 IEEE international conference on robotics and automation*. IEEE, 4290–4296.
- [45] Akshay Rangesch and Mohan Manubhai Trivedi. 2019. No blind spots: Full-surround multi-object tracking for autonomous vehicles using cameras and lidars. *IEEE Transactions on Intelligent Vehicles* 4, 4 (2019), 588–599.
- [46] John N Sanders-Reed and Phillip L Koon. 2010. Vision systems for manned and robotic ground vehicles. In *Unmanned Systems Technology XII*, Vol. 7692. SPIE, 181–192.
- [47] Ulrich Schwesinger, Mathias Bürki, Julian Timpner, Stephan Rottmann, Lars Wolf, Lina Maria Paz, Hugo Grimmett, Ingmar Posner, Paul Newman, Christian Häne, et al. 2016. Automated valet parking and charging for e-mobility. In *2016 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 157–164.
- [48] Ahmed Rida Sekkat, Yohan Dupuis, Varun Ravi Kumar, Hazem Rashed, Senthil Yogamani, Pascal Vasseur, and Paul Honeine. 2022. SynWoodScape: Synthetic surround-view fisheye camera dataset for autonomous driving. *IEEE Robotics and Automation Letters* 7, 3 (2022), 8502–8509.
- [49] Hyeonjun Sim, Jihyong Oh, and Munchul Kim. 2021. Xvfi: extreme video frame interpolation. In *Proceedings of the IEEE/CVF international conference on computer vision*. 14489–14498.
- [50] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [51] Dae-Young Song, Geonsoo Lee, HeeKyung Lee, Gi-Mun Um, and Donghyeon Cho. 2022. Weakly-supervised stitching network for real-world panoramic image generation. In *European Conference on Computer Vision*. Springer, 54–71.
- [52] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. 2018. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8934–8943.
- [53] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [54] Tian-Zhu Xiang, Gui-Song Xia, Xiang Bai, and Liangpei Zhang. 2018. Image stitching by line-guided local warping with global similarity constraint. *Pattern recognition* 83 (2018), 481–497.

1045	[55]	Jianxiong Xiao, Krista A Ehinger, Aude Oliva, and Antonio Torralba. 2012. Recognizing scene viewpoint using panoramic place representation. In <i>2012 IEEE Conference on Computer Vision and Pattern Recognition</i> . IEEE, 2695–2702.	1103
1046			1104
1047	[56]	Chunmei Xie, Xiaoyun Zhang, Hua Yang, Li Chen, and Zhiyong Gao. 2018. Video stitching based on optical flow. In <i>2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)</i> . IEEE, 1–5.	1105
1048			1106
1049	[57]	Weihua Ye, Kaiwen Yu, Yang Yu, and Jia Li. 2018. Logical stitching: A panoramic image stitching method based on color calibration box. In <i>2018 14th IEEE International Conference on Signal Processing (ICSP)</i> . IEEE, 1139–1143.	1107
1050			1108
1051	[58]	Senthil Yogamani, Ciarán Hughes, Jonathan Horgan, Ganesh Sistu, Padraig Varley, Derek O’Dea, Michal Uricár, Stefan Milz, Martin Simon, Karl Amende, et al. 2019. Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving. In <i>Proceedings of the IEEE/CVF International Conference on Computer Vision</i> . 9308–9318.	1109
1052			1110
1053			1111
1054			1112
1055			1113
1056			1114
1057			1115
1058			1116
1059			1117
1060			1118
1061			1119
1062			1120
1063			1121
1064			1122
1065			1123
1066			1124
1067			1125
1068			1126
1069			1127
1070			1128
1071			1129
1072			1130
1073			1131
1074			1132
1075			1133
1076			1134
1077			1135
1078			1136
1079			1137
1080			1138
1081			1139
1082			1140
1083			1141
1084			1142
1085			1143
1086			1144
1087			1145
1088			1146
1089			1147
1090			1148
1091			1149
1092			1150
1093			1151
1094			1152
1095			1153
1096			1154
1097			1155
1098			1156
1099			1157
1100			1158
1101			1159
1102			1160