

VITAMIN-E: Visual Tracking And Mapping with Extremely Dense Feature Points

Masashi Yokozuka, Shuji Oishi, Thompson Simon, Atsuhiko Banno
Robot Innovation Research Center,

National Institute of Advanced Industrial Science and Technology (AIST), Japan
{yokotsuka-masashi, shuji.oishi, simon.thompson, atsuhiko.banno}@aist.go.jp

Abstract

In this paper, we propose a novel indirect monocular SLAM algorithm called “VITAMIN-E,” which is highly accurate and robust as a result of tracking extremely dense feature points. Typical indirect methods have difficulty in reconstructing dense geometry because of their careful feature point selection for accurate matching. Unlike conventional methods, the proposed method processes an enormous number of feature points by tracking the local extrema of curvature informed by dominant flow estimation. Because this may lead to high computational cost during bundle adjustment, we propose a novel optimization technique, the “subspace Gauss–Newton method”, that significantly improves the computational efficiency of bundle adjustment by partially updating the variables. We concurrently generate meshes from the reconstructed points and merge them for an entire 3D model. The experimental results on the SLAM benchmark dataset EuRoC demonstrated that the proposed method outperformed state-of-the-art SLAM methods, such as DSO, ORB-SLAM, and LSD-SLAM, both in terms of accuracy and robustness in trajectory estimation. The proposed method simultaneously generated significantly detailed 3D geometry from the dense feature points in real time using only a CPU.

1. Introduction

Simultaneous localization and mapping (SLAM) is a key technology for applications such as autonomous systems and augmented reality. Whereas LiDAR-based SLAM[34, 11, 25] is well established and widely used in autonomous vehicles, visual SLAM with a monocular camera does not provide sufficient accuracy and robustness, particularly regarding dense map reconstruction, to replace LiDAR-based SLAM. Although some visual SLAM algorithms that use stereo cameras[7, 32], RGB-D cameras[12, 14, 33], and inertial sensors[20, 18, 2, 28, 19] have achieved high perfor-

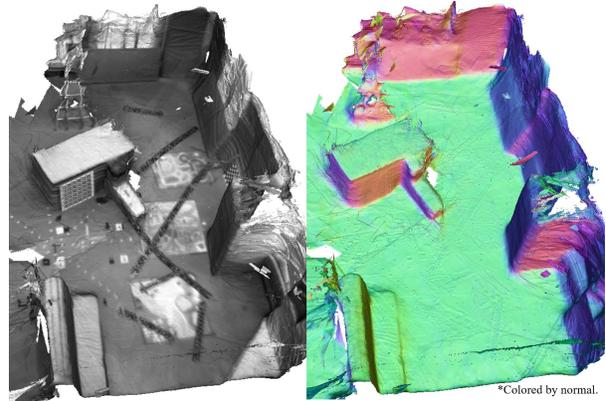


Figure 1. Dense geometry reconstruction with VITAMIN-E on EuRoC V101. (<https://youtu.be/yfKccCmmMsM>)

mance, these methods are based on pure monocular SLAM; hence, improving monocular SLAM is important.

Monocular SLAM methods can be classified into two types: direct methods and indirect methods.

Direct methods: Direct methods estimate camera poses and reconstruct the scene by minimizing the photometric error defined as a sum of the intensity difference between each pixel in the latest image and the reprojection of the color / monochrome 3D map. Direct methods, such as LSD-SLAM[5], SVO[9], and DSO[6], process almost all pixels in incoming images. They do not require exact pixel correspondences among multiple views unlike indirect methods, which leads to denser map reconstruction. However, direct methods are susceptible to image noise, luminance fluctuation, and lens aberration because they directly use pixel intensities. To overcome this drawback, Bergmann et al.[1] proposed a normalization method against luminance fluctuation and a calibration method for lens aberration. As another approach, Zhang et al.[35] proposed an auto-exposure method that is suitable for direct methods.

Indirect methods: Indirect methods minimize the geometric error between observed 2D feature points and re-

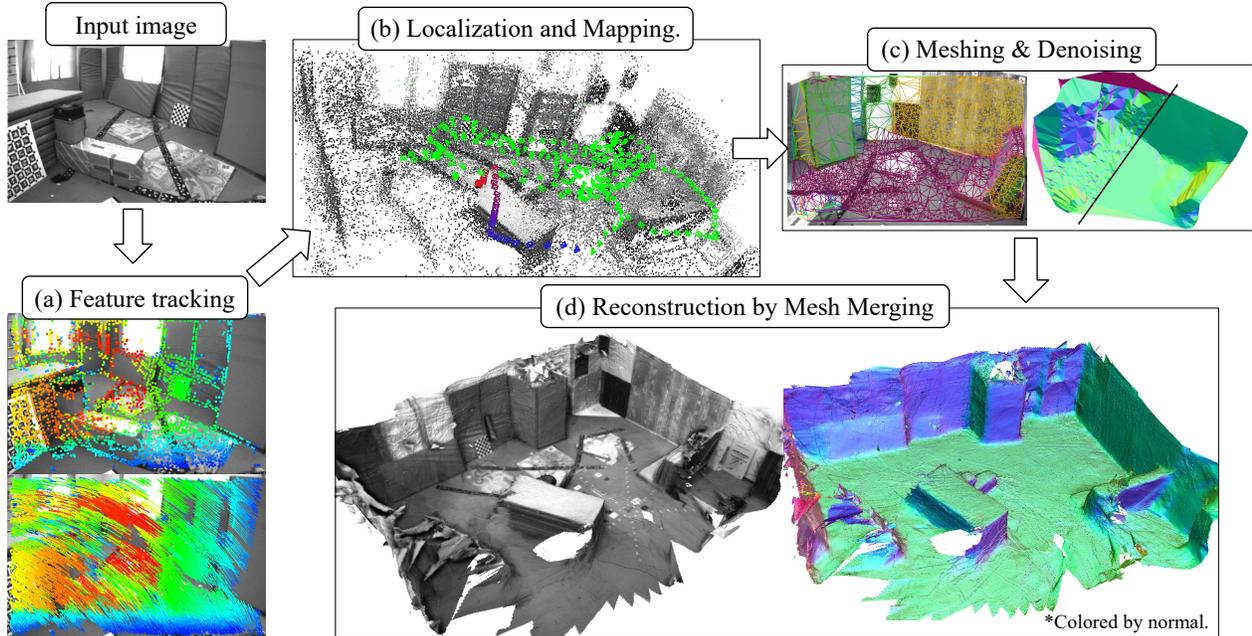


Figure 2. Overview of the proposed monocular SLAM.

projections of the corresponding 3D points. As a result of the use of feature descriptors, indirect methods such as PTAM[15] and ORB-SLAM[21] are robust against brightness changes and image noise. Additionally, indirect methods explicitly establish feature point correspondences; hence, outliers are easily removed using RANSAC[8] or M-estimation[27]. This characteristic, however, can be a drawback: Indirect methods carefully select stable feature points, thus the reconstructed 3D map tends to be sparse and does not provide detailed geometry. Densification methods, such as PMVS[10] and the extension L-PMVS[26], might be useful for obtaining dense geometry; however, they are offline methods and not applicable in real time.

In this paper, we propose the novel Visual Tracking And Mapping with Extremely dense feature points, “VITAMIN-E,” which is highly precise, robust, and dense because of the tracking of a large number of feature points. Indirect methods are inherently robust against noise, illumination change, and outliers as a result of the use of feature descriptors. Retaining this advantage, we reconstruct detailed 3D maps by establishing dense point correspondence. The contributions of this study are as follows: We first introduce a new dense feature point tracking algorithm based on dominant flow estimation and curvature extrema tracing. This allows VITAMIN-E to process an enormous number of feature points; however, the need to maintain them simultaneously might lead to a high computational cost. Therefore, we also introduce a novel optimization technique, called subspace Gauss–Newton method, for bundle adjustment. The optimization technique significantly improves

the efficiency of bundle adjustment by partially updating the variables. Moreover, VITAMIN-E generates meshes from the reconstructed feature points and integrates them using a truncated signed distance function (TSDF)[22, 30, 24]. Compared with not only conventional indirect methods but also state-of-the-art direct methods, VITAMIN-E provides highly detailed 3D geometry as shown in Figure 1 in real time using only a CPU

2. Dense Feature Point Tracking

2.1. Feature Point Tracking

Indirect methods that use image descriptors can be unstable because of incorrect feature point correspondences. They build feature point correspondences between multiple views by matching the descriptors. Extracting consistent descriptors over multiple frames, however, becomes difficult because descriptors vary as the monocular camera changes its pose. Methods such as the KLT tracker[29] that continuously track feature points while updating the feature descriptors might be useful for overcoming the problem. However, because the tracked positions drift as a result of a minute change of feature descriptors, the correspondences over multiple views tend to be incorrect. These problems originate with the use of feature descriptors.

Rather than associating feature points based on descriptors, VITAMIN-E tracks the local extrema of curvature in incoming images. In the proposed method, feature points denote the extrema of curvature on image intensities. Let $f(x, y)$ be an image, then curvature κ of image f is as fol-

lows:

$$\kappa = f_y^2 f_{xx} - 2f_x f_y f_{xy} + f_x^2 f_{yy}, \quad (1)$$

where f_x represents the partial derivative of f with respect to x , which can be obtained using a Sobel operator or similar technique. VITAMIN-E builds point correspondences over multiple images by tracking the local maximum point of curvature $\kappa(x, y, t)$, which is the extension of κ to time domain t . Figure 2(a) shows an example scene from which a large number of extrema of curvature κ are extracted. Whereas conventional indirect methods rely only on feature points with a large curvature to obtain stable correspondences, the proposed method tracks all detected extrema to reconstruct detailed geometry.

2.2. Dominant Flow Estimation

After detecting the extrema of curvature, the proposed method estimates a dominant flow that represents the average of optical flow over the images, which provides a good initial value to extrema tracking and makes it significantly stable, as explained later. Specifically, we determine the corresponding feature pairs between current and previous images using the BRIEF[4] feature. Because we only have to identify coarse feature pairs over consecutive frames at this moment, feature matching is performed on low-resolution images, subsampled to 1/6 of the former size.

Then, we fit the affine transformation model $\mathbf{y} = A\mathbf{x} + \mathbf{b}$ to the feature pairs. \mathbf{x} and \mathbf{y} denote the position of a feature point in the previous and current frame, respectively, and A and \mathbf{b} represent a matrix of 2×2 and a 2D translation, respectively. A and \mathbf{b} are obtained by minimizing cost function E using the Gauss–Newton method:

$$E = \sum_i^N \rho(\|\mathbf{y}_i - (A\mathbf{x}_i + \mathbf{b})\|_2), \quad (2)$$

where N and ρ denote the total number of corresponding points and a kernel function for M-estimation, respectively. The following Geman–McClure kernel with scale parameter σ is used in VITAMIN-E:

$$\rho(x) = \frac{x^2}{x^2 + \sigma^2}. \quad (3)$$

As a result of M-estimation, the dominant flow represented by A and \mathbf{b} can be estimated stably, even for low-resolution images, and it allows us to roughly predict the position of feature points in the next frame. Note that VITAMIN-E does not rely on conventional feature matching in its core but only for prior information for dense extrema tracking, as described in the next section. Whereas conventional feature matching has difficulty in making all feature points couple correctly between consecutive frames, affine transformation is easily obtained when at least three correspondences are given.

2.3. Curvature Extrema Tracking

VITAMIN-E tracks feature points by tracing the extrema of image curvature by making use of the dominant flow. Because it depends only on extrema instead of feature descriptors used in conventional indirect methods, VITAMIN-E is free from the variation of feature descriptors caused by image noise or illumination changes, which makes VITAMIN-E highly robust.

According to the dominant flow represented by A and \mathbf{b} , we first predict a current position $\bar{\mathbf{x}}_{t_1}$ of tracking point \mathbf{x}_{t_0} :

$$\bar{\mathbf{x}}_{t_1} = A\mathbf{x}_{t_0} + \mathbf{b}. \quad (4)$$

Next, prediction $\bar{\mathbf{x}}_{t_1}$ is corrected to \mathbf{x}_{t_1} by maximizing evaluation function F :

$$F = \kappa(\mathbf{x}_{t_1}, t_1) + \lambda w(\|\mathbf{x}_{t_1} - \bar{\mathbf{x}}_{t_1}\|_2), \quad (5)$$

where κ stores the curvature in each pixel, and $w(x) = 1 - \rho(x)$ and λ denote an evaluation function and weight for the prediction, respectively. The maximization is performed using the hill climbing method, with $\bar{\mathbf{x}}_{t_1}$ as the initial position. Specifically, maximum point \mathbf{x}_{t_1} is obtained by iterating the hill climbing method in eight neighboring pixels at each step until it reaches the local maximum value of F . Function w prevents the maximization process from falling into wrong extrema, thereby playing a regularization role.

Note that extrema tracking can easily fall into local solutions because there are many extrema in image curvature and it is almost impossible to distinguish them without any descriptors. However, the prediction according to the dominant flow boosts the accuracy of extrema tracking and enables it to approach the optimal solution.

3. Bundle Adjustment for Dense Tracking

3.1. Bundle Adjustment

Bundle adjustment iteratively adjusts the reconstructed map by minimizing reprojection errors. Given the i -th 3D point \mathbf{p}_i , the j -th camera's rotation R_j and translation \mathbf{t}_j , and the 2D position \mathbf{u}_{ij} of \mathbf{p}_i observed in the j -th camera frame, the objective function is formulated as follows:

$$E = \sum_i^N \sum_j^M \rho(\|\mathbf{u}_{ij} - \phi(R_j^T(\mathbf{p}_i - \mathbf{t}_j))\|_2), \quad (6)$$

where N and M are the numbers of feature points and camera poses respectively, ϕ denotes the 3D-2D projection function, and ρ is a kernel function for M-estimation. Specifically, optimal camera variables $\mathbf{c}_j = (R_j, \mathbf{t}_j)$ and \mathbf{p}_i are obtained by applying the Gauss–Newton method to Equation 6, which results in iteratively solving the following equations:

$$H\delta\mathbf{x} = -\mathbf{g}, \quad \mathbf{x} = \mathbf{x} + \delta\mathbf{x}, \quad (7)$$

where $\mathbf{x} = (\mathbf{c}_1, \dots, \mathbf{c}_M, \mathbf{p}_1, \dots, \mathbf{p}_N)^T$, and H and \mathbf{g} represent the Hessian matrix and gradient around \mathbf{x} of E , respectively. H and \mathbf{g} can be represented by the camera variable \mathbf{c}_j block and the feature point variable \mathbf{p}_i block as follows:

$$H = \begin{bmatrix} H_{cc} & H_{cp} \\ H_{cp}^T & H_{pp} \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} \mathbf{g}_c \\ \mathbf{g}_p \end{bmatrix}. \quad (8)$$

Hessian matrix H in bundle adjustment has a unique structure: H_{cc} and H_{pp} are sparse matrices with only diagonal elements in block units, whereas H_{cp} is a dense matrix. Efficient solutions that focus on this unique structure are the keys to developing highly precise and robust visual SLAM.

State-of-the-art monocular SLAM methods, such as ORB-SLAM[21] and DSO[6], solve Equation 7 by decomposing it using the Schur complement matrix instead of directly solving it:

$$\begin{aligned} (H_{cc} - H_{cp}H_{pp}^{-1}H_{cp}^T) \delta \mathbf{x}_c &= -\mathbf{g}_c + H_{cp}H_{pp}^{-1}\mathbf{g}_p, \quad (9) \\ H_{pp} \delta \mathbf{x}_p &= -\mathbf{g}_p - H_{cp}^T \delta \mathbf{x}_c, \quad (10) \end{aligned}$$

where $\mathbf{x}_c = (\mathbf{c}_1, \dots, \mathbf{c}_M)^T$ and $\mathbf{x}_p = (\mathbf{p}_1, \dots, \mathbf{p}_N)^T$. The decomposition allows us to solve bundle adjustment faster. The number of camera variables M is remarkably smaller than that of feature point variables, and the inverse matrix of H_{pp} can be easily calculated because it has only diagonal components; thus, the Schur complement matrix $(H_{cc} - H_{cp}H_{pp}^{-1}H_{cp}^T)$ whose size is $M \times M$ is significantly tiny compared with the original H , and the inverse matrix is rapidly computable.

Equation 9 is also called marginalization. When regarding $H_{cc} - H_{cp}H_{pp}^{-1}H_{cp}^T$ as a new H and $\mathbf{g}_c - H_{cp}H_{pp}^{-1}\mathbf{g}_p$ as a new \mathbf{g} , the decomposition is equivalent to eliminating all feature point variables p from cost function E . State-of-the-art SLAMs make themselves efficient using the marginalization technique to prevent the increase in computational cost caused by a large number of variables.

However, in the case of maintaining thousands of feature points in every frame, as in the dense extrema tracking in VITAMIN-E, the size of matrix H fundamentally cannot be made sufficiently small because variable elimination is applicable only to old variables unrelated to the current frame for stability. Moreover, the size of the Schur complement matrix is proportional to the number of feature points; thus, the calculation cost of bundle adjustment over tens of thousands points, where the size of H is $100,000 \times 100,000$ or more, becomes too high to run bundle adjustment in real time.

3.2. Subspace Gauss–Newton Method

To deal with the explosion in the size of H , we propose a novel optimization technique called the “subspace Gauss–Newton method.” It partially updates variables rather than

updating all of them at once, as in Equations 9 and 10, by decomposing these equations further as follows:

$$H_{c_i c_i} \delta \mathbf{x}_{c_i} = - \left(\mathbf{g}_{c_i} + \sum_{l=1}^{i-1} H_{c_l c_i} \delta \mathbf{x}_{c_l} + \sum_{r=i+1}^M H_{c_i c_r} \delta \mathbf{x}_{c_r} + \sum_{j=1}^N H_{c_i p_j} \delta \mathbf{x}_{p_j} \right), \quad (11)$$

$$H_{p_j p_j} \delta \mathbf{x}_{p_j} = - \left(\mathbf{g}_{p_j} + \sum_{l=1}^{j-1} H_{p_l p_j} \delta \mathbf{x}_{p_l} + \sum_{r=j+1}^N H_{p_i p_r} \delta \mathbf{x}_{p_r} + \sum_{i=1}^M H_{c_i p_j}^T \delta \mathbf{x}_{c_i} \right). \quad (12)$$

Equation 11 updates $\delta \mathbf{x}_{c_i}$ of a camera variable, and Equation 12 $\delta \mathbf{x}_{p_j}$ of a feature point variable. $H_{c_i c_i}$, $H_{c_i p_j}$, and $H_{p_j p_j}$ are matrices of 6×6 , 6×3 , and 3×3 , respectively. The subspace Gauss–Newton method iteratively solves Equations 11 and 12 until $\delta \mathbf{x}_{c_i}$ and $\delta \mathbf{x}_{p_j}$ converge, respectively.

These formulae are extensions of the Gauss–Seidel method, which is an iterative method to solve a linear system of equations, and equivalent to solving the Gauss–Newton method by fixing all variables except the variables to be optimized. The advantage of the proposed subspace Gauss–Newton method is that it does not require a large inverse matrix, unlike Equation 9, but instead, only an inverse matrix of 6×6 at most. Additionally, as in ORB-SLAM[21] and DSO[6], further speedup is possible by appropriately performing variable elimination that sets most elements of H_{cc} and H_{pp} to zero. Because the proposed optimization method limits the search space in the Gauss–Newton method to its subspace, we call it the “subspace Gauss–Newton method.”¹

4. Dense Reconstruction

A large number of accurate 3D points are generated in real time with VITAMIN-E as a result of the dense extrema tracking and subspace Gauss–Newton method described in previous sections. This leads not only to point cloud generation but also allows further dense geometry reconstruction that cannot be achieved by conventional indirect methods.

Meshing and Noise Removal: We first project the 3D points onto an image and apply Delaunay triangulation to generate triangular meshes. Then, We use NLTGV minimization proposed by Greene et al[23] to remove noise on the meshes. NLTGV minimization allows us to smooth the meshes, thereby retaining local surface structures, unlike

¹ Alternating optimization, such as our method, has been used in some contexts[36][31]. See the supplementary information for details of the novel aspect of our method.

typical mesh denoising methods such as Laplacian smoothing. Figure 2(c) shows example results of Delaunay triangulation and smoothing with NLTGV minimization.

Mesh Integration in TSDF: Finally, we integrate the meshes in a TSDF to reconstruct the entire geometry of the scene. The TSDF represents an object shape by discretizing the space into grids that store the distance from the object surface, and can merge multiple triangular meshes by storing the average value of distances from the meshes to each grid.

5. Experimental Results

5.1. Setup

We evaluated the performance of the proposed VITAMIN-E on the visual SLAM benchmark EuRoC[3].² The dataset was created using flying drones equipped with a stereo camera and an IMU in an indoor environment, and provided ground truth of trajectories obtained by a Leica MS 50 laser tracker and Vicon motion capture. EuRoC is also well-known for data variations, with different difficulties ranked by the movement speed and lighting conditions. In this experiment, we compared results obtained using VITAMIN-E and other monocular SLAM methods, DSO[6], ORB-SLAM[21], and LSD-SLAM[5], using only the left images of the stereo camera in the EuRoC dataset. Note that because VITAMIN-E and DSO[6] do not include loop closing and relocalization, these functions in ORB-SLAM[21] and LSD-SLAM[5] were disabled to evaluate performance fairly. Similar evaluations can be found in the papers on DSO[6] and SVO[9].

VITAMIN-E ran on a Core i7-7820 HQ without any GPUs, threading each process for real time processing. Initialization was performed using essential matrix decomposition. Bundle adjustment is significantly sensitive to the initial value of the variable. In VITAMIN-E, we initialized the camera variables using P3P[17, 13] with RANSAC[8] and feature point variables using triangulation. Note that the proposed bundle adjustment ran so fast that we applied it to every frame rather than each key frame as in ORB-SLAM and DSO. To manage the TSDF, OpenChisel[16] was used in our implementation because of its ability to handle a TSDF on a CPU.

5.2. Evaluation Criteria

The evaluation on EuRoC was performed according to the following criteria:

Localization success rate: We defined the localization success rate as $\frac{N_s}{N}$, where N_s and N denote the number of images that were successfully tracked and the images in the entire sequence, respectively. If the success rate was less

² See the supplementary information for experimental results on other datasets.

than 90%, we regarded the trial as a failure. When localization failed even once, the methods could not estimate the camera position later because loop closing and relocalization were disabled in this experiment. Therefore, robustness greatly contributed to the success rate.

Localization accuracy: The localization accuracy was computed by scaling the estimated trajectories so that the RMSE from ground truth trajectories was minimized because scale is not available in monocular SLAM. Note that we did not evaluate the accuracy when the success rate was less than 90% because the RMSE of very short trajectories tends to have an unfairly high accuracy.

Number of initialization retries: Initialization plays an important role in monocular visual SLAM and significantly affects the success rate. Because different methods have different initialization processes, the number of initialization retries in each method is not directly comparable, but can be a reference regarding whether the method has a weakness in initialization in certain cases.

5.3. Results and Discussion

Table 1 shows the experimental results. The results were obtained by applying each method to the image sequences in EuRoC five times, and Table 1 shows the average values and the standard deviations of the aforementioned criteria. Bold font is used to emphasize the highest accuracy in each sequence. Regarding LSD-SLAM, the results of initialization retries are not included in the table because LSD-SLAM did not have a re-initialization function and failed to initialize in any sequences. We thus manually identified the frame for which the initialization worked well in each trial, so the number of retries of LSD-SLAM was excluded from the evaluation.

The EuRoC image sequences MH01, MH02, MH04, MH05, V101, V201, and V202 are relatively easy cases for visual SLAM because the camera motion is relatively slow and the illumination does not change frequently. By contrast, the camera moves fast in MH03, V102, V103, and V203, and additionally the lighting conditions dynamically change in V102, V103, and V203. Furthermore, in V203, the exposure time is so long that we can see severe motion blur in the image sequence, particularly in an extremely dark environment.

Even for the challenging environments, VITAMIN-E never lost localization and outperformed other SLAM methods, DSO, ORB-SLAM, and LSD-SLAM, both in terms of accuracy and robustness, as shown in Table 1. Particularly, in the sequences that contain fast camera motion, such as V102 and V103, the proposed method was superior to the existing methods. The reconstruction results are shown in Figures 3 and 4. Despite the proposed VITAMIN-E successfully generating dense and accurate geometry compared with its competitors, it performed equally fast on a CPU, as

Table 1. Experimental results : localization success or failure [\checkmark or \times], localization accuracy [cm], localization success rate [%], and number of initialization retries.

Sequence name (no. of images)	Our method			DSO[6]			ORB-SLAM[21] w/o loop closure			LSD-SLAM[5] w/o loop closure		
MH01 easy (3682)	\checkmark	12.9 ± 0.5 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	6.0 ± 0.8 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	5.2 ± 1.1 97.7 ± 1.6 19 ± 11	cm % %	\times	(44.9 ± 7.2) 28.9 ± 23.6 -	cm % %
MH02 easy (3040)	\checkmark	8.8 ± 0.5 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	4.2 ± 0.2 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	4.1 ± 0.4 92.4 ± 1.1 56 ± 6	cm % %	\times	(58.3 ± 6.9) 73.0 ± 1.5 -	cm % %
MH03 medium (2700)	\checkmark	10.6 ± 1.3 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	21.1 ± 0.9 100.0 ± 0.0 0 ± 0	cm % %	\times	(4.5 ± 0.4) 48.9 ± 0.8 0 ± 0	cm % %	\times	(266.2 ± 61.3) 28.4 ± 20.7 -	cm % %
MH04 difficult (2033)	\checkmark	19.3 ± 1.6 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	20.3 ± 1.0 95.7 ± 0.0 5 ± 0	cm % %	\checkmark	33.6 ± 9.4 95.2 ± 0.8 6 ± 1	cm % %	\times	(136.4 ± 114.3) 27.2 ± 7.0 -	cm % %
MH05 difficult (2273)	\checkmark	14.7 ± 1.1 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	10.2 ± 0.6 95.5 ± 0.0 2 ± 0	cm % %	\checkmark	14.9 ± 4.6 90.0 ± 4.0 18 ± 5	cm % %	\times	(27.4 ± 16.4) 22.7 ± 0.5 -	cm % %
V101 easy (2911)	\checkmark	9.7 ± 0.2 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	13.4 ± 5.8 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	8.8 ± 0.1 96.6 ± 0.0 1 ± 0	cm % %	\times	(20.0 ± 22.8) 11.6 ± 11.2 -	cm % %
V102 medium (1710)	\checkmark	9.3 ± 0.6 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	53.0 ± 5.5 100.0 ± 0.0 0 ± 0	cm % %	\times	(14.5 ± 11.7) 52.0 ± 3.3 17 ± 4	cm % %	\times	(67.0 ± 14.0) 15.2 ± 0.1 -	cm % %
V103 difficult (2149)	\checkmark	11.3 ± 0.5 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	85.0 ± 36.4 100.0 ± 0.0 0 ± 0	cm % %	\times	(37.2 ± 20.7) 65.5 ± 8.8 56 ± 26	cm % %	\times	(29.3 ± 2.0) 11.0 ± 0.1 -	cm % %
V201 easy (2280)	\checkmark	7.5 ± 0.4 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	7.6 ± 0.5 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	6.0 ± 0.1 95.2 ± 0.0 0 ± 0	cm % %	\times	(131.3 ± 20.4) 74.1 ± 8.9 -	cm % %
V202 medium (2348)	\checkmark	8.6 ± 0.7 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	11.8 ± 1.4 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	12.3 ± 2.7 99.5 ± 1.2 0 ± 0	cm % %	\times	(42.1 ± 9.2) 11.3 ± 0.2 -	cm % %
V203 difficult (1922)	\checkmark	14.0 ± 5.2 100.0 ± 0.0 0 ± 0	cm % %	\checkmark	147.5 ± 6.6 100.0 ± 0.0 0 ± 0	cm % %	\times	(104.3 ± 64.0) 16.8 ± 15.9 233 ± 123	cm % %	\times	(17.7 ± 1.6) 11.9 ± 0.2 -	cm % %

Table 2. Average tracking time per frame in each method.

Our method	DSO	ORB-SLAM	LSD-SLAM
36 msec/frame	53 msec/frame	25 msec/frame	30 msec/frame

Table 3. Average computation time for each process of VITAMIN-E. Whereas the front-end processes ran in parallel for each frame, the back-end processes for generating the 3D mesh model were performed at a certain interval.

Front-end			Back-end	
Feature tracking	Localization & mapping	Meshing & denoising	TSDF updating & marching cubes (low-resolution; voxel size $\simeq 15$ cm)	TSDF updating & marching cubes (high-resolution; voxel size $\simeq 2.5$ cm)
36 msec/frame	25 msec/frame	45 msec/frame	175 msec/time	4000 msec/time

shown in Tables 2 and 3. Note that the smaller the size of the voxel in a TSDF for a detailed 3D model, the higher the computational cost.

The high accuracy and robustness of VITAMIN-E derives from tracking a large number of feature points and performing bundle adjustment for every frame. Sharing reprojection errors among variables is important for accurate monocular SLAM, and the proposed method efficiently diffuses errors to an enormous number of variables via fast bundle adjustment. Simultaneously, it prevents localization failure caused by losing sight of some feature points by han-

dling a large number of feature points.

6. Conclusion

In this paper, we proposed a novel visual SLAM method that reconstructs dense geometry with a monocular camera. To process a large number of feature points, we proposed curvature extrema tracking using the dominant flow between consecutive frames. The subspace Gauss-Newton method was also introduced to maintain an enormous number of variables by partially updating them to avoid a large

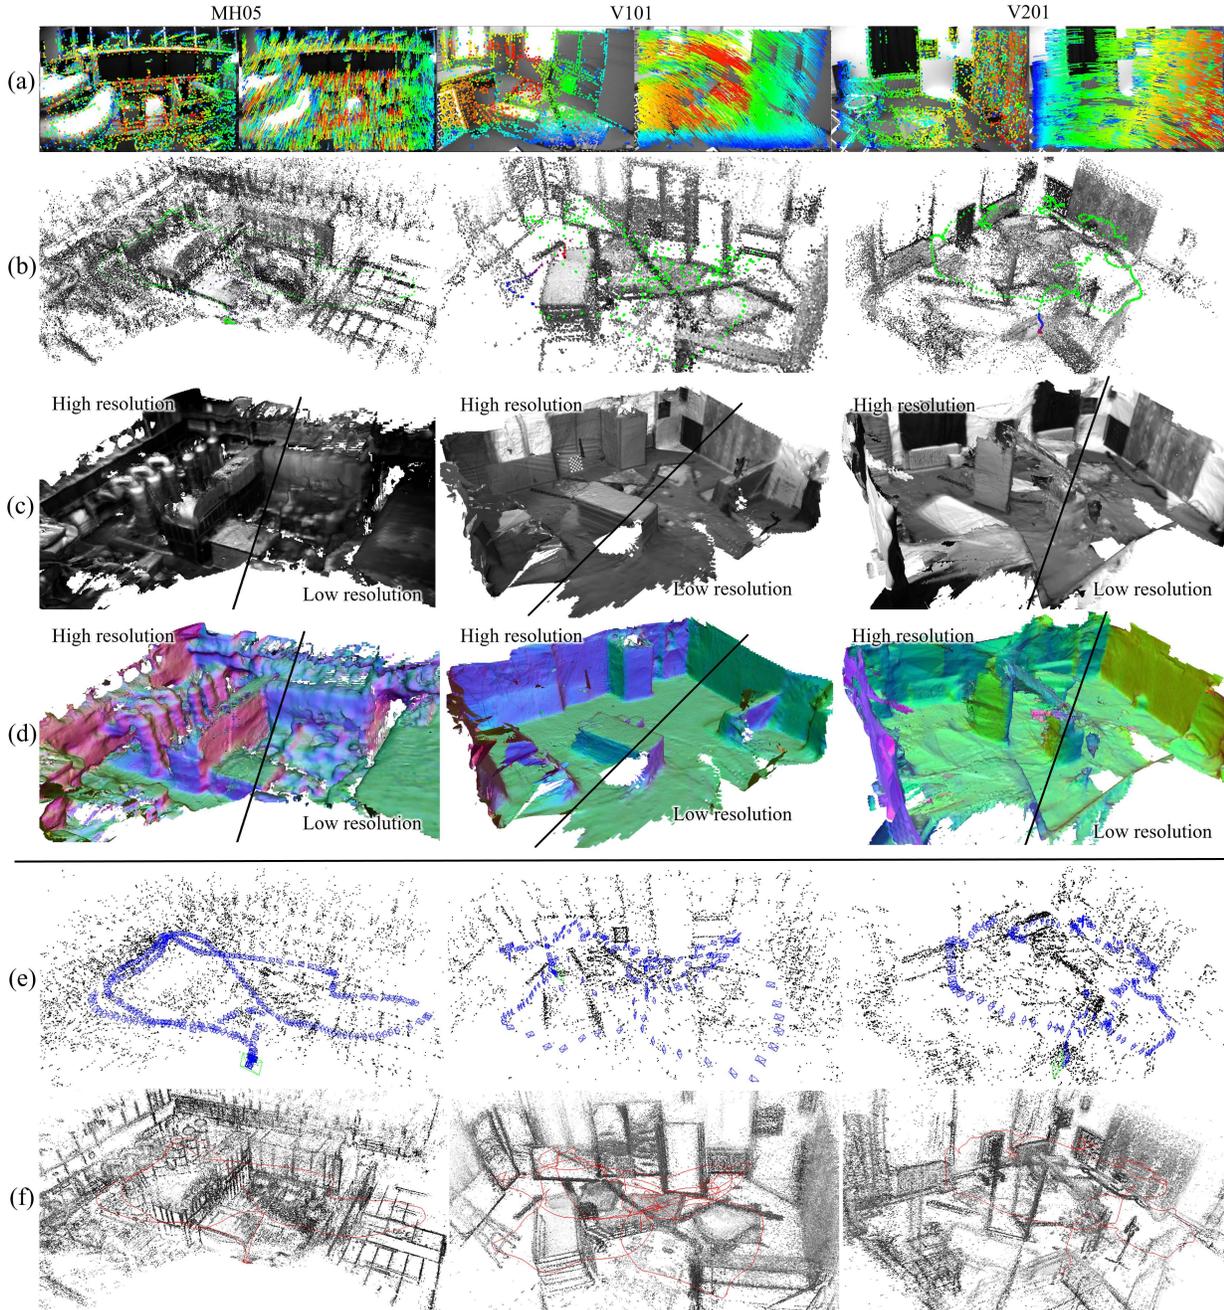


Figure 3. Reconstruction results : (a) dense extrema tracking in real time, (b) reconstructed 3D points, (c) mesh models, and (d) normal maps generated with the proposed dense geometry reconstruction in different sized TSDF voxels, and reconstructed 3D points in the same scenes with (e) ORB-SLAM and (f) DSO.

inverse matrix calculation in bundle adjustment. Moreover, supported by the accurate and dense point clouds, we achieved highly dense geometry reconstruction with NLTGV minimization and TSDF integration.

VITAMIN-E is executable on a CPU in real time, and it outperformed state-of-the-art SLAM methods, DSO, ORB-SLAM, and LSD-SLAM, both in terms of accuracy and robustness on the EuRoC benchmark dataset. Performance

should be improved when loop closing is introduced in VITAMIN-E, and fusing IMU data would also be effective to stably estimate the camera pose for challenging environments such as EuRoC V203.

Acknowledgements: This work was supported in part by JSPS KAKENHI, Japan Grant Numbers 16K16084 and 18K18072, and a project commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

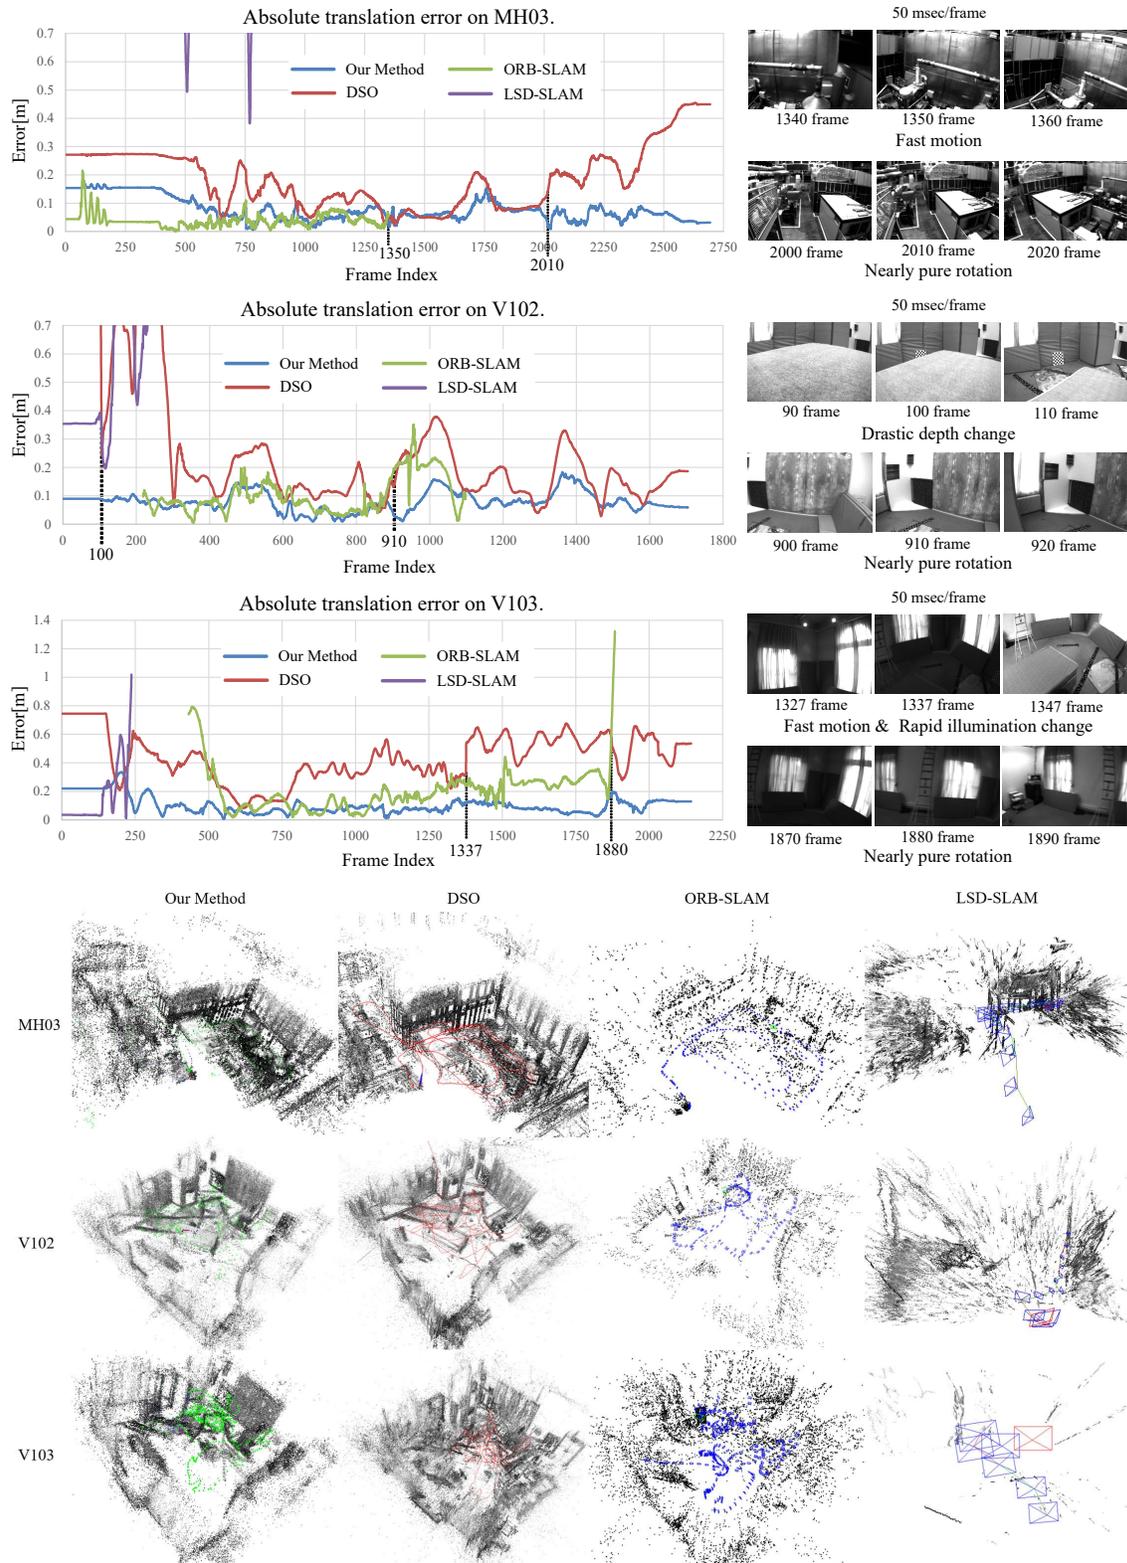


Figure 4. Chronological RMSEs of estimated trajectories on EuRoC V102, V103, and MH03, and the reconstructed point clouds in each scene: VITAMIN-E successfully estimated the camera trajectories despite a drastic depth change, nearly pure camera rotation, rapid camera motion, and severe lighting conditions, whereas the competitors suffered from them and resulted in large trajectory errors or completely getting lost.

References

- [1] Paul Bergmann, Rui Wang, and Daniel Cremers. Online photometric calibration of auto exposure video for realtime visual odometry and SLAM. *IEEE Robotics and Automation Letters*, 3, 2018.
- [2] Michael Blosch, Sammy Omari, Marco Hutter, and Roland Siegwart. Robust visual inertial odometry using a direct ekf-based approach. In *Proc. of International Conference on Intelligent Robots (IROS)*, 2015.
- [3] Michael Burri, Janosch Nikolic, Pascal Gohl, Thomas Schneider, Joern Rehder, Sammy Omari, Markus W Achtelik, and Roland Siegwart. The EuRoC micro aerial vehicle datasets. *International Journal of Robotics Research*, 35(10):1157–1163, 2016.
- [4] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. BRIEF: Binary robust independent elementary features. In *Proc. of European Conference on Computer Vision (ECCV)*, 2010.
- [5] Jakob Engel and Daniel Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In *Proc. of European Conference on Computer Vision (ECCV)*, 2014.
- [6] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [7] Jakob Engel, Jörg Stückler, and Daniel Cremers. Large-scale direct slam with stereo cameras. In *Proc. of International Conference on Intelligent Robots (IROS)*, 2015.
- [8] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [9] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. SVO: Semidirect visual odometry for monocular and multicamera systems. *IEEE Transactions on Robotics*, 33(2):249–265, 2017.
- [10] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multi-view stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8):1362–1376, 2010.
- [11] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2D LIDAR SLAM. In *Proc. of International Conference on Robotics and Automation (ICRA)*, 2016.
- [12] Albert S. Huang, Abraham Bachrach, Peter Henry, Michael Krainin, Daniel Maturana, Dieter Fox, and Nicholas Roy. Visual odometry and mapping for autonomous flight using an RGB-D camera. In *Proc. of International Symposium on Robotics Research (ISRR)*, 2011.
- [13] Tong Ke and Stergios Roumeliotis. An efficient algebraic solution to the perspective-three-point problem. In *Proc. of International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [14] Christian Kerl, Jürgen Sturm, and Daniel Cremers. Robust odometry estimation for RGB-D cameras. In *Proc. of International Conference on Robotics and Automation (ICRA)*, 2013.
- [15] Georg Klein and David Murray. Parallel tracking and mapping for smaller workspaces. In *Proc. of International Symposium on Mixed and Augmented Reality (ISMAR)*, 2007.
- [16] Matthew Klingensmith, Ivan Dryanovski, Siddhartha Srinivasa, and Jizhong Xiao. Chisel: Real time large scale 3D reconstruction onboard a mobile device using spatially hashed signed distance fields. *Robotics: Science and Systems*, 2015.
- [17] Laurent Kneip, Davide Scaramuzza, and Roland Siegwart. A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In *Proc. of International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [18] Stefan Leutenegger, Simon Lynen, Michael Bosse, Roland Siegwart, and Paul Furgale. Keyframe-based visual-inertial odometry using nonlinear optimization. *International Journal of Robotics Research*, 34(3):314–334, 2015.
- [19] Simon Lynen, Markus Wilhelm Achtelik, Stephan Weiss, and Margarita Chli. A robust and modular multi-sensor fusion approach applied to mav navigation. In *Proc. of International Conference on Intelligent Robots (IROS)*, 2013.
- [20] Anastasios I. Mourikis and Stergios I. Roumeliotis. A multi-state constraint Kalman filter for vision-aided inertial navigation. In *Proc. of International Conference on Robotics and Automation (ICRA)*, 2007.
- [21] Raul Mur-Artal and Juan D. Tardós. ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, 2017.
- [22] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *Proc. of International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011.
- [23] W. Nicholas Greene and Nicholas Roy. Flame: Fast lightweight mesh estimation using variational smoothing on delaunay graphs. In *Proc. of International Conference on Computer Vision (ICCV)*, 2017.
- [24] Helen Oleynikova, Zachary Taylor, Marius Fehr, Roland Siegwart, and Juan Nieto. Voxblox: Incremental 3D euclidean signed distance fields for on-board mav planning. In *Proc. of International Conference on Intelligent Robots (IROS)*, 2017.
- [25] Chanoh Park, Peyman Moghadam, Sooshin Kim, Alberto Elfes, Clinton Fookes, and Sridha Sridharan. Elastic lidar fusion: Dense map-centric continuous-time slam. In *Proc. of International Conference on Robotics and Automation (ICRA)*, 2018.
- [26] Alex Poms, Chenglei Wu, Shou-I Yu, and Yaser Sheikh. Learning patch reconstructability for accelerating multi-view stereo. In *Proc. of International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [27] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, USA, second edition, 1992.
- [28] Tong Qin, Peiliang Li, and Shaojie Shen. Vins-mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Transactions on Robotics*, 34(4):1004–1020, 2018.

- [29] Jianbo Shi and Carlo Tomasi. Good features to track. In *Proc. of International Conference on Computer Vision and Pattern Recognition (CVPR)*, 1994.
- [30] Frank Steinbrücker, Jürgen Sturm, and Daniel Cremers. Volumetric 3D mapping in real-time on a cpu. In *Proc. of International Conference on Robotics and Automation (ICRA)*, 2014.
- [31] Bill Triggs, Philip McLauchlan, Richard Hartley, and Andrew Fitzgibbon. Bundle adjustment - a modern synthesis. In *Proc. of the Int. WS on Vision Algorithms: Theory and Practice*, 2000.
- [32] R. Wang, M. Schwörer, and D. Cremers. Stereo DSO: Large-scale direct sparse visual odometry with stereo cameras. In *International Conference on Computer Vision (ICCV)*, 2017.
- [33] Ji Zhang, Michael Kaess, and Sanjiv Singh. Real-time depth enhanced monocular odometry. In *Proc. of International Conference on Intelligent Robots (IROS)*, 2014.
- [34] Ji Zhang and Sanjiv Singh. LOAM: Lidar odometry and mapping in real-time. *Robotics: Science and Systems*, 2014.
- [35] Zichao Zhang, Christian Forster, and Davide Scaramuzza. Active exposure control for robust visual odometry in HDR environments. In *Proc. of International Conference on Robotics and Automation (ICRA)*, 2017.
- [36] Qian-Yi Zhou and Vladlen Koltun. Color map optimization for 3d reconstruction with consumer depth cameras. *ACM Trans. Graph.*, 33(4):155:1–155:10, 2014.