Asynchronous Multi-View SLAM

Anqi Joyce Yang[*,1,2], Can Cui[*,1,3], Ioan Andrei Bârsan[*,1,2], Raquel Urtasun[1,2], Shenlong Wang[1,2]

Abstract— **Existing multi-camera SLAM systems assume synchronized shutters for all cameras, which is often not the case in practice. In this work, we propose a generalized multi-camera SLAM formulation which accounts for asynchronous sensor observations. Our framework integrates a continuous-time motion model to relate information across asynchronous multi-frames during tracking, local mapping, and loop closing. For evaluation, we collected *AMV-Bench*, a challenging new SLAM dataset covering 482 km of driving recorded using our asynchronous multi-camera robotic platform. AMV-Bench is over an order of magnitude larger than previous multi-view HD outdoor SLAM datasets, and covers diverse and challenging motions and environments. Our experiments emphasize the necessity of asynchronous sensor modeling, and show that the use of multiple cameras is critical towards robust and accurate SLAM in challenging outdoor scenes. For additional information, please see the project website at:**
https://www.cs.toronto.edu/~ajyang/amv-slam

I. INTRODUCTION

Simultaneous Localization and Mapping (SLAM) is the task of localizing an autonomous agent in unseen environments by building a map at the same time. SLAM is a fundamental part of many technologies ranging from augmented reality to photogrammetry and robotics. Due to the availability of camera sensors and the rich information they provide, camera-based SLAM, or visual SLAM, has been widely studied and applied in robot navigation.

Existing visual SLAM methods [1]–[5] and benchmarks [6]–[8] mainly focus on either monocular or stereo camera settings. Although lightweight, such configurations are prone to tracking failures caused by occlusion, dynamic objects, lighting changes and textureless scenes, all of which are common in the real world. Many of these challenges can be attributed to the narrow field of view typically used (Fig. 1a). Due to their larger field of view (Fig. 1b), wide-angle or fisheye lenses [9], [10] or multi-camera rigs [11]–[16] can significantly increase the robustness of visual SLAM systems [15].

Nevertheless, using multiple cameras comes with its own set of challenges. Existing stereo [5] or multi-camera [11]–[15] SLAM literature assumes synchronized shutters for all cameras and adopts discrete-time trajectory modeling based on this assumption. However, in practice different cameras are not always triggered at the same time, either due to technical limitations, or by design. For instance, the camera shutters could be synchronized to another sensor, such as a spinning LiDAR (*e.g.*, Fig. 1c), which is a common set-up in self-driving [17]–[20]. Moreover, failure to account for the robot motion in between the firing of the cameras could lead to localization failures. Consider a car driving along a highway at 30m/s (108km/h). Then in a short 33ms camera firing interval, the vehicle would travel one meter, which is significant when centimeter-accurate pose estimation is required. As a result, a need arises for a generalization of multi-view visual SLAM to be agnostic to camera timing, while being scalable and robust to real-world conditions.

In this paper we formalize the *asynchronous multi-view SLAM (AMV-SLAM)* problem. Our first contribution is a general framework for AMV-SLAM, which, to the best of our knowledge, is the first full asynchronous continuous-time multi-camera visual SLAM system for large-scale outdoor environments. Key to this formulation is (1) the concept of asynchronous multi-frames, which group input images from multiple asynchronous cameras, and (2) the integration of a continuous-time motion model, which relates spatio-temporal information from asynchronous multi-frames for joint continuous-time trajectory estimation.

Since there is no public asynchronous multi-camera SLAM dataset, our second contribution is *AMV-Bench*, a novel large-scale dataset with high-quality ground-truth. AMV-Bench was collected during a full year in Pittsburgh, PA, and includes challenging conditions such as low-light scenes, occlusions, fast driving (Fig. 1d), and complex maneuvers like three-point turns and reverse parking. Our experiments show that multi-camera configurations are critical in overcoming adverse conditions in large-scale outdoor scenes. In addition, we show that asynchronous sensor modeling is crucial, as treating the cameras as synchronous leads to 30% higher failure rate and $4\times$ the local pose errors compared to asynchronous modeling.

II. RELATED WORK

1) Visual SLAM / Visual Odometry: SLAM has been a core area of research in robotics since the 1980s [21]–[25]. The comprehensive survey by Cadena et al. [26] provides a detailed overview of SLAM. Modern visual SLAM approaches can be divided into direct and indirect methods. Direct methods like DTAM [27], LSD-SLAM [1], and DSO [3] estimate motion and map parameters by directly optimizing over pixel intensities (photometric error) [28], [29]. Alternatively, indirect methods, which are the focus of this work, minimize the re-projection energy (geometric error) [30] over an intermediate representation obtained from raw images. A common subset of these are feature-based methods like PTAM [31] and ORB-SLAM [4] which represent raw observations as sets of keypoints.

(a) FoV of a stereo pair (b) FoV of 5 wide-angle cameras

0 ms	33.3 ms	66.7ms	100 ms	...
Stereo pair	Wide rear right	Wide rear left	Stereo pair	
Wide front middle		Wide front left	Wide front middle	
Wide front right			Wide front right	

(c) Asynchronous camera firing timeline in our dataset

| Stereo left | Stereo Right | Front Middle | Front Right | Rear Right | Rear Left | Front Left |

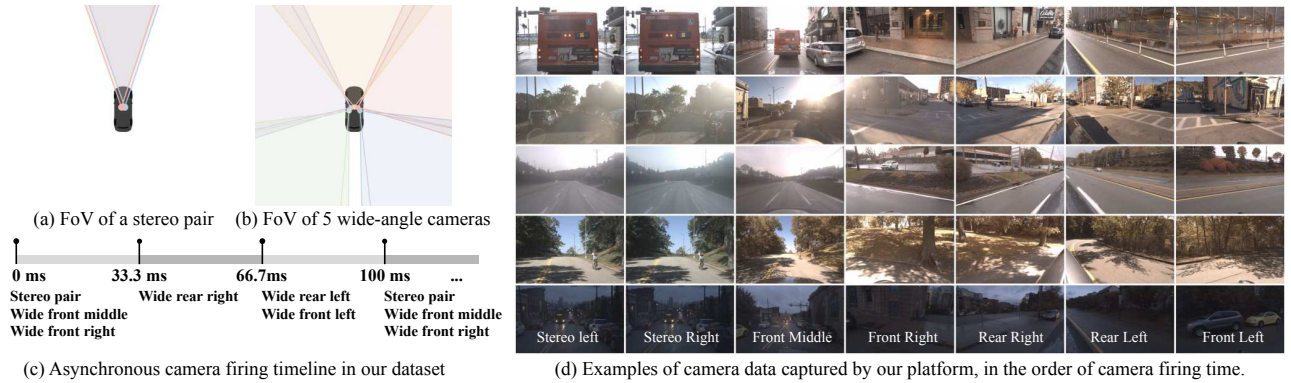(d) Examples of camera data captured by our platform, in the order of camera firing time.

Fig. 1: The asynchronous multi-camera rig in AMV-Bench, containing a stereo pair and five wide-angle cameras. The cameras are synced to a LiDAR, with the asynchronous firing schedule shown in (c). The sample images highlight challenging scenarios like occlusions, sunlight glare, low-textured highways, shadows on the road, and low-light rainy environments.

2) Multi-View SLAM: Monocular [1], [3], [4], [31] and stereo [2], [5], [32] are the most common visual SLAM configurations. However, many applications could benefit from a much wider field of view for better perception and situation awareness during navigation. Several multi-view SLAM approaches have been proposed [12]–[15], [33]–[38]. Early filtering-based approaches [33] treat multiple cameras as independent sensors, and fuse their observations using a central Extended Kalman Filter. Recent optimization-based multi-camera SLAM systems [12]–[15] extend monocular PTAM [31], ORB-SLAM [4] and DSO [3] respectively to synchronized multi-camera rigs to jointly estimate ego-poses at discrete timestamps. Multi-view visual-inertial systems such as VINS-MKF [36] and ROVINS [38] also assume synchronous camera timings.

3) Continuous-time Motion Models: Continuous-time motion models help relate sensors triggered at arbitrary times while moving, with applications such as calibration [39], [40], target tracking [41], [42] and motion planning [43], [44]. Continuous-time SLAM typically focuses on visual-inertial fusion [45], [46], rolling-shutter cameras [46]–[51] or LiDARs [52]–[55]. Klingner et al. [56] propose a continuous-time Structure-from-Motion framework for multiple synchronous rolling-shutter cameras. The key component for continuous-time trajectory modeling is choosing a family of functions that is both flexible and reflective of the kinematics. A common approach is fitting parametric functions over the states, *e.g.*, piecewise linear functions [50], [56], spirals [57], wavelets [58], or B-splines [45], [59], [60]. Other approaches represent trajectories through non-parametric methods such as Gaussian Processes [39], [40], [55], [61]–[63].

III. NOTATION

1) Coordinate Frame: We denote a coordinate frame x with $\mathcal{F}_x$. $\mathbf{T}_{yx}$ is the rigid transformation that maps homogeneous points from $\mathcal{F}_x$ to $\mathcal{F}_y$. In this work we use three coordinate frames: the world frame $\mathcal{F}_w$, the moving robot's body frame $\mathcal{F}_b$, and the camera frame $\mathcal{F}_k$ associated with each camera $\mathcal{C}_k$.

2) Pose and Motion: The pose of a 3D rigid body can be represented as a rigid transform from $\mathcal{F}_b$ to $\mathcal{F}_w$ as follows:

$$\mathbf{T}_{wb} = \begin{bmatrix} \mathbf{R}_{wb} & \mathbf{t}_w \\ \mathbf{0}^T & 1 \end{bmatrix} \in \mathbb{SE}(3) \quad \text{with} \quad \mathbf{R}_{wb} \in \mathbb{SO}(3), \mathbf{t}_w \in \mathbb{R}^3$$

where $\mathbf{R}_{wb}$ is the 3×3 rotation matrix, $\mathbf{t}_w$ is the translation vector, and $\mathbb{SE}(3)$ and $\mathbb{SO}(3)$ are the Special Euclidean and Special Orthogonal Matrix Lie Groups respectively. We define the trajectory of a 3D rigid body as a function $\mathbf{T}_{wb}(t) : \mathbb{R} \to \mathbb{SE}(3)$ over the time domain $t \in \mathbb{R}$.

3) Lie Algebra Representation: For optimization purposes, a 6-DoF minimal pose representation associated with the Lie Algebra $\mathfrak{se}(3)$ of the matrix group $\mathbb{SE}(3)$ is widely adopted. It is a vector $\boldsymbol{\xi} = [\, \boldsymbol{v}^T \; \boldsymbol{\omega}^T \,]^T \in \mathbb{R}^6$, where $\boldsymbol{v} \in \mathbb{R}^3$ and $\boldsymbol{\omega} \in \mathbb{R}^3$ encode the translation and rotation components respectively. We use the uppercase Exp (and, conversely, Log) to convert $\boldsymbol{\xi} \in \mathbb{R}^6$ to $\mathbf{T} \in \mathbb{SE}(3)$: $\text{Exp}(\boldsymbol{\xi}) = \exp(\boldsymbol{\xi}^\wedge) = \mathbf{T}$, where $\exp$ is the matrix exponential, and $\boldsymbol{\xi}^\wedge = \begin{bmatrix} \boldsymbol{\omega}_\times & \boldsymbol{v} \\ \mathbf{0}^T & 0 \end{bmatrix} \in \mathfrak{se}(3)$ with $\boldsymbol{\omega}_\times$ being the 3×3 skew-symmetric matrix of $\boldsymbol{\omega}$.

4) Motion Model: We use superscripts to denote the type of motion models. c and ℓ represent the cubic B-spline motion model $\mathbf{T}^c(t)$ and the linear motion model $\mathbf{T}^\ell(t)$ respectively.

IV. ASYNCHRONOUS MULTI-VIEW SLAM

We consider the asynchronous multi-view SLAM problem where the observations are captured by multiple cameras triggered at arbitrary times with respect to each other. Each camera $\mathcal{C}_k$ is assumed to be a calibrated pinhole camera with intrinsic matrix $\mathbf{K}_k$, and extrinsics encoded by the mapping $\mathbf{T}_{kb}$ from the body frame $\mathcal{F}_b$ to camera frame $\mathcal{F}_k$. The input to the problem is a sequence of image and capture timestamp pairs $\{(I_{ik}, t_{ik})\}_{\forall i}$, associated with each camera $\mathcal{C}_k$. The goal is then to estimate the robot trajectory $\mathbf{T}_{wb}(t)$ in the world frame. As a byproduct we also estimate a map $\mathcal{M}$ of the 3D structure of the environment as a set of points.

Our system follows the standard visual SLAM structure of initialization coupled with the three-threaded tracking, local mapping, and loop closing, with the key difference that we generalize to multiple cameras with asynchronous timing via *asynchronous multi-frames* (Sec. IV-A) and a
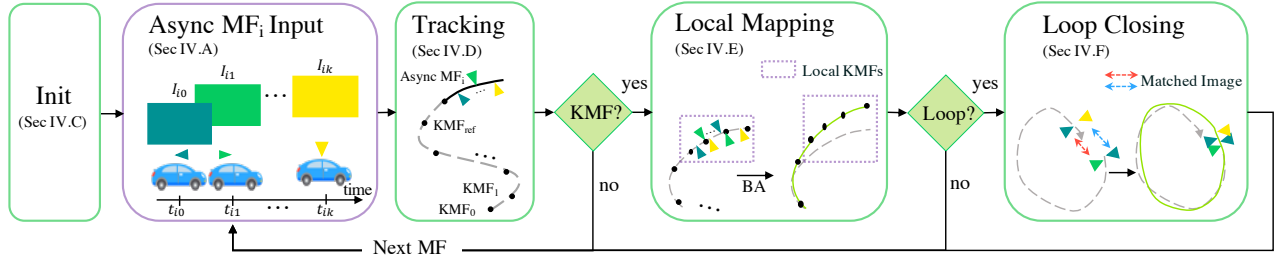
2

Fig. 2: Overview. Initialization is followed by tracking, local mapping and loop closing. MF=Multi-Frame; KMF=Key MF.
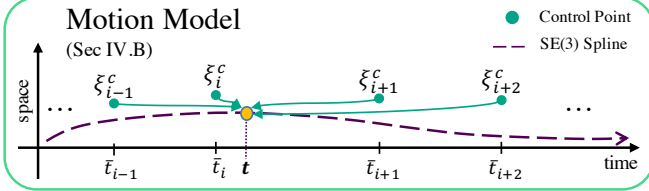

Fig. 3: Illustration of the cubic B-spline model. The pose $\mathbf{T}_{wb}^c(t)$ at time $t \in [\bar{t}_i, \bar{t}_{i+1}]$ is defined by four control poses associated with key multi-frames indexed at $i-1, i, i+1, i+2$.

continuous-time motion model (Sec. IV-B). In particular, after initialization (Sec. IV-C), tracking (Sec. IV-D) takes each incoming multi-frame as input, infers its motion parameters, and decides whether to promote it as a *key* multi-frame (KMF). For efficiency, only KMFs are used during local mapping (Sec. IV-E) and loop closing (Sec. IV-F). When a new KMF is selected, the local mapping module refines poses and map points over a recent window of KMFs to ensure local consistency, while the loop closing module detects when a previously-mapped area is being revisited and corrects the drift to enhance global consistency. See Fig. 2 for an overview.

A. Asynchronous Multi-Frames

Existing synchronous multi-view systems [14] group multi-view images captured at the same time into a *multi-frame* as input. However, this cannot be directly applied when the firing time of each sensor varies. To generalize to asynchronous camera timings, we introduce the concept of *asynchronous multi-frame*, which groups images that are captured closely (*e.g.*, within 100ms) in time. In Fig. 1 each asynchronous multi-frame contains the images taken during a single spinning LiDAR sweep at 10 Hz. Contrasting to synchronous multi-frames [14] that store images and a discrete pose estimated at a single timestamp, each *asynchronous* multi-frame MF_i stores: (1) a set of image and capture time pairs $\{(I_{ik}, t_{ik})\}$ indexed by associated camera $\mathcal{C}_k$, and (2) continuous-time motion model parameters to recover the estimated trajectory.

B. Continuous-Time Trajectory Representation

To associate the robot pose with observations that could be made at arbitrary times, we formulate the overall robot trajectory as a *continuous-time* function $\mathbf{T}_{wb}(t) : \mathbb{R} \to \mathbb{SE}(3)$, rather than discrete poses. We exploit a *cumulative* cubic B-spline function [45] as the first and second derivatives of this parameterization are smooth and computationally efficient to evaluate [59]. The cumulative structure is necessary

for accurate on-manifold interpolation in $\mathbb{SE}(3)$ [45], [64]. Following [45], given a knot vector $\mathbf{b} = [t_0, t_1, \ldots, t_7] \in \mathbb{R}^8$, a cumulative cubic B-spline trajectory $\mathbf{T}_{wb}^c(t)$ over $t \in [t_3, t_4)$ is defined by four control poses $\boldsymbol{\xi}_0^c : \boldsymbol{\xi}_3^c \in \mathbb{R}^6$ [45]. In our framework, we associate each key multi-frame KMF_i with a control pose. In addition, since the KMFs do not necessarily distribute evenly in time, we use a non-uniform knot vector. For each KMF_i, we define the *representative time* $\bar{t}_i$ as the median of all image capture times t_{ik}, and define the knot vector as $\mathbf{b}_i = [\bar{t}_{i-3}, \bar{t}_{i-2}, \ldots, \bar{t}_{i+4}] \in \mathbb{R}^8$. Then, the spline trajectory over the interval $t \in [\bar{t}_i, \bar{t}_{i+1})$ can be expressed as a function of four control poses $\boldsymbol{\xi}_{i-1}^c, \boldsymbol{\xi}_i^c, \boldsymbol{\xi}_{i+1}^c, \boldsymbol{\xi}_{i+2}^c$:

$$\mathbf{T}_{wb}^c(t) = \text{Exp}\left(\boldsymbol{\xi}_{i-1}^c\right) \prod_{j=1}^{3} \text{Exp}\left(\widetilde{B}_{j,4}(t)\boldsymbol{\Omega}_{i-1+j}\right), \quad (1)$$

where $\boldsymbol{\Omega}_{i-1+j} = \text{Log}(\text{Exp}(\boldsymbol{\xi}_{i-2+j}^c)^{-1}\text{Exp}(\boldsymbol{\xi}_{i-1+j}^c))$ is the relative pose between control poses, and $\widetilde{B}_{j,4}(t) = \sum_{l=j}^{3} B_{l,4}(t) \in \mathbb{R}$ is the cumulative basis function, where the basis function $B_{l,4}(t)$ is computed with the knot vector $\mathbf{b}_i$ using the de Boor-Cox formula [65], [66]. See Fig. 3 for an illustration, and the appendix for more details.

C. Initialization

The system initialization assumes that there exists a pair of cameras that share a reasonable overlapping field of view and fire very closely in time (*e.g.*, a synchronous stereo pair, present in most autonomous vehicle setups [19], [67], [68]). At the system startup, we create the first MF with the associated camera images and capture times, select it as the first KMF, set the representative time $\bar{t}_0$ to the camera pair firing time, the control pose $\boldsymbol{\xi}_0^c$ to the origin of the world frame, and initialize the map with points triangulated from the camera pair. Map points from other camera images are created during mapping after the second KMF is inserted.

D. Tracking

During tracking, we estimate the continuous pose of an incoming multi-frame MF_i by matching it with the most recent KMF. We then decide whether MF_i should be selected as a KMF for map refinement and future tracking. Following [4], we formulate pose estimation and map refinement as an indirect geometric energy minimization problem based on sparse image features.

1) Feature Matching: For each image in the new MF, we identify its reference images in the reference KMF as images captured by the same camera or any camera sharing a reasonable overlapping field of view. We extract sparse 2D keypoints and associated descriptors from the new images and match them against the reference image keypoints to establish associations with existing 3D map points. We denote the set of matches as $\{(\mathbf{u}_{i,k,j}, \mathbf{X}_j)\}_{\forall(k,j)}$, where $\mathbf{u}_{i,k,j} \in \mathbb{R}^2$ is the 2D keypoint extracted from image I_{ik} in MF_i, and $\mathbf{X}_j \in \mathbb{R}^3$ is the matched 3D map point in the world frame.

2) Pose Estimation: Cubic B-splines are effective for modeling the overall trajectory, but directly using them in tracking entails estimating four 6-DoF control poses that define motion not only in the new MF, but also in the existing trajectory. Therefore, more information is needed for a stable estimation. For computational efficiency, we instead use a simpler and less expressive continuous-time linear motion model $\mathbf{T}_{wb}^\ell(t)$ during tracking, whose parameters are later used to initialize the cubic B-spline model in Sec. IV-E. Specifically, we estimate MF_i pose $\boldsymbol{\xi}_i^\ell \in \mathbb{R}^6$ at the representative timestamp $\bar{t}_i$, and evaluate the continuous pose with linear interpolation and extrapolation: $\mathbf{T}_{wb}^\ell(t) = \text{Exp}(\boldsymbol{\xi}_i^\ell)\left(\text{Exp}(\boldsymbol{\xi}_i^\ell)^{-1}\mathbf{T}_{wb}^c(\bar{t}_{\text{ref}})\right)^\alpha$, where $\alpha = \frac{\bar{t}_i - t}{\bar{t}_i - \bar{t}_{\text{ref}}}$ and $\bar{t}_{\text{ref}}$ is the representative timestamp of the reference KMF.

Coupled with the obtained multi-view correspondences, we formulate pose estimation for $\boldsymbol{\xi}_i^\ell$ as a constrained, asynchronous, multi-view case of the perspective-n-points (PnP) problem, in which a geometric energy is minimized:

$$E_{geo}(\boldsymbol{\xi}_i^\ell) = \sum_{(k,j)} \rho\left(\left\|\mathbf{e}_{i,k,j}(\boldsymbol{\xi}_i^\ell)\right\|_{\boldsymbol{\Sigma}_{i,k,j}^{-1}}^2\right), \qquad (2)$$

where $\mathbf{e}_{i,k,j}(\boldsymbol{\xi}_i^\ell) \in \mathbb{R}^2$ is the reprojection error of the matched correspondence pair $(\mathbf{u}_{i,k,j}, \mathbf{X}_j)$, and $\boldsymbol{\Sigma}_{i,k,j} \in \mathbb{R}^{2\times2}$ is a diagonal covariance matrix denoting the uncertainty of the match. ρ denotes a robust norm, with Huber loss used in practice. The reprojection error for the pair is defined as:

$$\mathbf{e}_{i,k,j}(\boldsymbol{\xi}_i^\ell) = \mathbf{u}_{i,k,j} - \pi_k\left(\mathbf{X}_j, \mathbf{T}_{wb}^\ell(t_{ik})\mathbf{T}_{kb}^{-1}\right), \qquad (3)$$

where $\pi_k(\cdot, \cdot) : \mathbb{R}^3 \times \mathbb{SE}(3) \rightarrow \mathbb{R}^2$ is the perspective projection function of camera $\mathcal{C}_k$, $\mathbf{T}_{kb}$ is the respective camera extrinsics matrix, and $\mathbf{T}_{wb}^\ell(t)$ is the linear model used only during tracking to initialize the cubic B-spline model later.

We initialize $\boldsymbol{\xi}_i^\ell$ by linearly extrapolating poses from the previous two multi-frames MF_{i-2} and MF_{i-1} based on a constant-velocity model. To achieve robustness against outlier map point associations, we wrap the above optimization in a RANSAC loop, where only a minimal number of $(\mathbf{u}_{i,k,j}, \mathbf{X}_j)$ are sampled to obtain each hypothesis. We solve the optimization with the Levenberg-Marquardt (LM) algorithm within each RANSAC iteration. Given our initialization, the problem converges in a few steps in each RANSAC iteration.

3) Key Multi-Frame Selection: We use a hybrid key frame selection scheme based on estimated motion [31] and map point reobservability [4]. Namely, the current MF is registered as a KMF if the tracked pose has a local translational or rotational change above a certain threshold, or if the ratio of map points reobserved in a number of cameras is below a certain threshold. In addition, to better condition the shape of cubic B-splines, a new KMF is regularly inserted during long periods of little change in motion and scenery (*e.g.*, when the robot stops). Empirically we find reobservability-only heuristics insufficient during very fast motions in low-textured areas (*e.g.*, fast highway driving), but show their combination with the motion-based heuristic to be robust to such scenarios.

E. Local Mapping

When a new KMF is selected, we run local bundle adjustment to refine the 3D map structure and minimize drift accumulated from tracking errors in recent frames. Map points are then created and culled to reflect the latest changes.

1) Bundle Adjustment: We use windowed bundle adjustment to refine poses and map points in a set of recent KMFs. Its formulation is similar to the pose estimation problem, except extended to a window of N key frames $\{\boldsymbol{\xi}_i^c\}_{1\leq i\leq N}$ to jointly estimate a set of control poses and map points:

$$E_{geo}(\{\boldsymbol{\xi}_i^c\}, \{\mathbf{X}_j\}) = \sum_{(i,k,j)} \rho\left(\left\|\mathbf{e}_{i,k,j}(\{\boldsymbol{\xi}_i^c\}, \mathbf{X}_j)\right\|_{\boldsymbol{\Sigma}_{i,k,j}^{-1}}^2\right). \quad (4)$$

Note that unlike tracking, we now refine the estimated local trajectory with the cubic B-spline model $\mathbf{T}_{wb}^c(t)$ parameterized by control poses $\{\boldsymbol{\xi}_i^c\}$. We initialize the control pose $\boldsymbol{\xi}_N^c$ of the newly-inserted key multi-frame with $\boldsymbol{\xi}_N^\ell$ estimated in tracking. For observations made after $\bar{t}_{N-1}$, their pose evaluation would involve control poses $\boldsymbol{\xi}_{N+1}^c$ and $\boldsymbol{\xi}_{N+2}^c$ and knot vector values $\bar{t}_{N+1\leq p\leq N+4}$ which do not yet exist. To handle such boundary cases, we represent these future control poses and timestamps as a linear extrapolation function of existing control poses and timestamps. We again minimize Eq. (4) with the LM algorithm.

2) Map Point Creation and Culling: With a newly-inserted KMF, we triangulate new map points with the refined poses and keypoint matches from overlapping image pairs both within the same KMF and across neighboring KMFs. To increase robustness against dynamic objects and outliers, we cull map points that are behind the cameras or have a reprojection errors above a certain threshold.

F. Loop Closing

The loop closing module detects when the robot revisits an area, and corrects the accumulated drift to achieve global consistency in mapping and trajectory estimation. With a wider field of view, multi-view SLAM systems can detect loops that are encountered at arbitrary angles.

We extend the previous DBoW3 [69]-based loop detection algorithm [4] with a multi-view similarity check and a multi-view geometric verification. To perform loop closure, we integrate the cubic B-spline motion model to formulate an asynchronous, multi-view case of the pose graph optimization problem. Please see the appendix for details.

V. DATASET

Much of the recent progress in computer vision and robotics has been driven by the existence of large-scale high-quality datasets [6], [70], [72], [73]. However, in the field

4

TABLE I: An overview of major SLAM datasets. Legend: W = diverse weather, G = large geographic diversity, MV = multi-view, HD = vertical resolution ≥ 1080. *total travel distance not explicitly released at the time of writing.

Name	Total km	Async	W	G	MV	HD
KITTI-360 [10]	74				✓	✓
RobotCar [70]	1000		✓		✓	
Ford Multi-AV [67]	n/A*		✓		✓	✓
A2D2 [68]	≈100*		✓	✓	✓	✓
4Seasons [32]	350		✓	✓		
EU Long-Term [71]	1000		✓		✓	✓
Ours	482	✓	✓	✓	✓	✓

TABLE II: Baseline methods. M=monocular, S=stereo, A=all cameras. RPE-T(cm/m), RPE-R(rad/m), ATE(m), AUC(%).

Method	RPE-T		RPE-R		ATE		SR(%)
	med	AUC	med	AUC	med	AUC	
DSO-M [75]	42.72	28.08	8.02E-05	54.23	594.39	44.67	62.67
ORB-M [4]	34.00	25.66	5.49E-05	63.77	694.37	42.65	64.00
ORB-S [5]	1.85	65.70	3.29E-05	70.47	30.74	74.31	77.33
Sync-S	1.30	77.54	2.91E-05	78.37	24.53	77.44	84.00
Sync-A [14]	2.15	68.46	3.47E-05	70.47	58.18	75.01	74.67
Ours-A	**0.35**	**88.63**	**1.13E-05**	**88.17**	**6.13**	**88.82**	**92.00**

of SLAM, in spite of their large number, previous datasets have been insufficient for evaluating asynchronous multi-view SLAM systems due to either scale, diversity, or sensor configuration limitations. Such datasets either emphasize a specific canonical route over a long period of time, foregoing geographic diversity [67], [70], [71], do not have a surround camera configuration critical for robustness [6], [32], [70], or lack the large scale necessary for evaluating safety-critical SLAM systems [6], [10], [68], [74]. Furthermore, none of the existing SLAM datasets feature multiple asynchronous cameras to directly evaluate an AMV-SLAM system.

To address these limitations, we propose *AMV-Bench*, a novel large-scale asynchronous multi-view SLAM dataset recorded using a fleet of SDVs in Pittsburgh, PA over the span of one year. Table I shows a high-level comparison between the proposed dataset and other similar SLAM-focused datasets. Please see the appendix for details.

1) Sensor Configuration: Each vehicle is equipped with seven cameras, wheel encoders, an IMU, a consumer-grade GPS receiver, and an HDL-64E LiDAR. The LiDAR data is only used to compute the ground-truth poses. There are five wide-angle cameras spanning most of the vehicle's surroundings, and an additional forward-facing stereo pair. All intrinsic and extrinsic calibration parameters are computed in advance using a set of checkerboard calibration patterns. All images are rectified to a pinhole camera model.

Each camera has an RGB resolution of 1920×1200 pixels, and uses a global shutter. The five wide-angle cameras are hardware triggered in sync with the rotating LiDAR at an average frequency of 10Hz, firing asynchronously with respect to each other. Fig. 1 illustrates the camera configuration.

2) Dataset Organization: The dataset contains 116 sequences spanning 482km and 21h. All sequences are recorded during daytime or twilight. Each sequence ranges from 4 to 18 minutes, with a wide variety of scenarios including (1) diverse environments (busy streets, highways, residential and rural areas) (2) diverse weather ranging from sunny days to heavy precipitation; (3) diverse motions with varying speed (highway, urban traffic, parking lots), trajectory loops, and maneuvers including U-turns and reversing. Please refer to Fig. 1d for examples.

The dataset is split geographically into train, validation, and test subsets (65/25/26 sequences), as shown in the appendix. The ground-truth poses are obtained using an offline HD map-based global optimization leveraging IMU, odometer,

GPS, and LiDAR. The maps are built from multiple runs through the same area, ensuring ground-truth accuracy.

VI. Experiments

We evaluate our method on the proposed AMV-Bench dataset. We first show that it outperforms several popular SLAM methods [4], [5], [14], [75]. Next, we perform ablation studies highlighting the importance of asynchronous modeling, the use of multiple cameras, the impact of loop closure and, finally, the differences between feature extractors.

A. Experimental Details

1) Implementation Details: All images are downsampled to 960×600 for both our method as well as all baselines. In our system, we extract 1000 ORB [76] keypoints from each image, using grid-based sampling [4] to encourage homogeneous distribution. Matching is performed with nearest neighbor + Lowe's ratio test [77] with threshold 0.7. A new KMF is inserted either (1) when the estimated local translation against reference KMF is over 1m, or local rotation is over 1°, or (2) when under 35% of the map points are re-observed in at least two camera frames, or (3) when a KMF hasn't been inserted for 20 consecutive MFs. (3) is necessary to model the spline trajectory when the robot stays stationary. We perform bundle adjustment over a recent window of size $N = 11$.

2) Experiment Set-Up & Metrics: We use the training set (65 sequences) for hyperparameter tuning, and the validation set (25 sequences) for testing. To account for stochasticity, we run each experiments three times. We use three classic metrics: absolute trajectory error (ATE) [78], relative pose error (RPE) [6], and success rate (SR). SR is the fraction of sequences that were successfully completed without SLAM failures such as tracking loss or repeated mapping failures. We evaluate ATE at 10Hz and RPE at 1Hz. For each method, we report the mean SR over the three trials.

For a large-scale dataset like AMV-Bench, it is impractical to list the ATE and RPE errors for each sequence. To aggregate results, for each method we report the median and the area under a cumulative error curve (AUC) of the errors in all trials at all evaluated timestamps. Missing entries due to SLAM failures are padded with infinity. AUC is computed between 0 and a given threshold, which we set to be 20cm/m, 5E-04rad/m and 1km for RPE-T, RPE-R and ATE respectively.

B. Results

1) Quantitative Comparison: We compare our method with multiple popular SLAM methods, including monocular

TABLE III: (Left) Ablation on motion models. (Right) Ablation on cameras, all with the cubic B-spline model; s=stereo, wf=wide-front, wb=wide-back. All ablations performed with loop closing disabled.

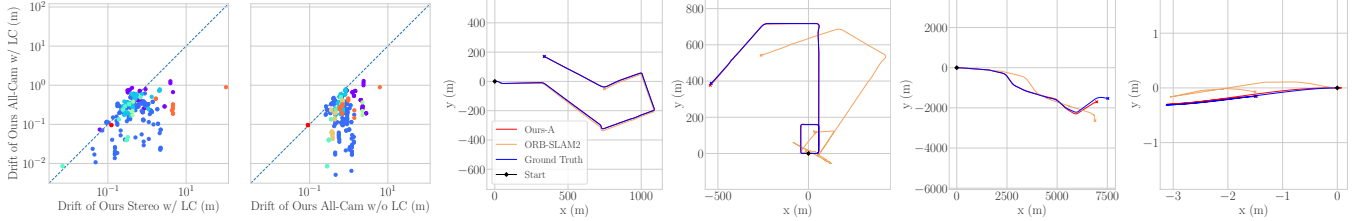|	RPE-T med	RPE-T AUC	RPE-R med	RPE-R AUC	ATE med	ATE AUC	SR	Cameras s	wf	wb	RPE-T med	RPE-T AUC	RPE-R med	RPE-R AUC	ATE med	ATE AUC	SR
Synch.	1.97	69.46	2.96E-05	73.39	55.24	75.11	70.7	✓			0.70	79.86	1.93E-05	80.48	11.44	75.75	88.0
Linear	0.41	87.76	**1.11E-05**	88.39	**6.09**	88.31	89.3	✓	✓		0.41	84.88	1.21E-05	85.86	9.00	84.92	90.7
B-spline	**0.35**	**88.79**	**1.11E-05**	**88.47**	6.53	**89.04**	**92.0**	✓	✓	✓	**0.35**	**88.79**	**1.11E-05**	**88.47**	**6.53**	**89.04**	**92.0**

Fig. 4: (Left) Pose drift (1) after multi-view/stereo loop closure; (2) with/without multi-view loop closure. Colors correspond to different sequences. (Right) Qualitative results. Rightmost is a maneuver reversing into a parallel parking spot.

DSO [3], [75], and monocular [4] and stereo [5] ORB-SLAM. We use the front middle camera in the monocular setting. We also compare with our discrete-time motion model implementation, using (1) only the stereo cameras, and (2) all 7 cameras, which corresponds to a multi-view sync baseline [14]. All methods are run with loop closure. Table V shows that our dataset is indeed challenging, with third-party baselines finishing under 80% of the validation sequences. Our method significantly outperforms the rest in terms of accuracy and robustness. Our stereo-sync baseline performs better than [5] mostly due to the more robust key frame selection strategy.

2) Motion Model Ablation: We run the following experiments using all cameras but with different motion models: (1) a synchronous discrete-time model; (2) an asynchronous linear motion model; (3) an asynchronous cubic B-spline model, which is our proposed solution. Loop closing is disabled for simplicity. Table III shows that the wrong synchronous timing assumption finishes about 30% fewer sequences and has local errors that are $4\times$ as high compared to the main system. Furthermore, trajectory modeling with cubic B-splines consistently outperforms the less expressive linear model.

3) Camera Ablation: We experiment with different camera combinations with the same underlying cubic B-spline motion model. We disable loop closing for simplicity. Table III indicates a performance boost in all metrics with more cameras covering a wider field of view.

4) Loop Closure: We first study the effect of multiple cameras on loop detection. Out of 9 validation sequences containing loops, our full method using all cameras could detect 8 loops with 100% precision, while our stereo loop detection implementation was only able to detect 6 loops closed in the same direction. The leftmost subfigure of Fig. 4 compares the drift after multi-view vs. stereo loop closure. To further highlight the reduction in global trajectory drift, the second subfigure in Fig. 4 depicts the drift relative to the ground truth with and without loop closure at every multi-frame where loop closure was performed.

5) Features: Indirect SLAM typically uses classic keypoint extractors [76], [77], yet recently, learning-based extractors [79]–[83] have shown promising results. We benchmark a set of classic [76], [84] and learned [79], [83] keypoint extractors. For learned methods, we directly run the provided pre-trained models without re-training. Loop closure is disabled for simplicity. Our results show that SIFT and SuperPoint finish more sequences than ORB (98.7%/98.7% SR vs. 92.0% for ORB) and have higher ATE AUC (91.5%/96.3% vs. 89.0%), with the caveat that feature extraction is much slower.

6) Runtime: Unlike monocular or stereo SLAM methods, asynchronous multi-view SLAM requires processing multiple camera images (seven in our setting). This introduces a linear multiplier to the complexity of the full SLAM pipeline, including feature extraction, tracking, bundle adjustment, and loop closure. Thus, despite the significant improvement on robustness and accuracy, our current implementation is not able to achieve real-time operation. Improving the efficiency of AMV-SLAM is thus an important area for future research.

7) Qualitative Results: Fig. 4 plots trajectories of our method, ORB-SLAM2 [5] and the ground truth in selected validation sequences. Our approach outperforms ORB-SLAM2 and visually aligns well with the GT trajectories in most cases. We also showcase a failure case from a rainy highway sequence. For additional quantitative and qualitative results, please see the appendix.

VII. Conclusion

In this paper, we formalized the problem of multi-camera SLAM with asynchronous shutters. Our framework groups input images into asynchronous multi-frames, and extends feature-based SLAM to the asynchronous multi-view setting using a cubic B-spline continuous-time motion model.

To evaluate AMV-SLAM systems, we proposed a new large-scale asynchronous multi-camera outdoor SLAM dataset, AMV-Bench. Experiments on this dataset highlight the necessity of the asynchronous sensor modeling, and the importance of using multiple cameras to achieve robustness and accuracy in challenging real-world conditions.
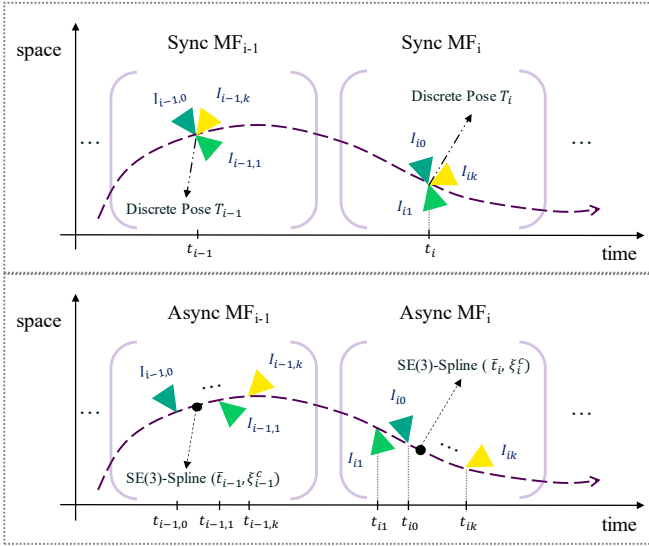
Fig. 5: Illustration of the synchronous multi-frame vs. the asynchronous multi-frame. For simplicity, in this illustration each async multi-frame is assumed to be a key multi-frame. Each key multi-frame is associated with cubic B-spline motion parameters to define the overall trajectory.

APPENDIX

A. OVERVIEW

The appendix covers additional information on:

1) Our method, specifically the details of the linear motion model used during tracking, more mathematical details of cubic B-splines, as well as further details on the loop closing module.
2) Our dataset and its geographic splits, providing a more in-depth comparison to other related SLAM benchmarks.
3) Our experiments, showcasing additional ablation studies, quantitative tables, qualitative results and discussions.

B. METHOD

A. Asynchronous Multi-Frames

We provide an illustration for the concept of an asynchronous multi-frame compared to a synchronous multi-frame in Fig. 5.

B. Linear Motion Model

During tracking, we estimate poses in the current asynchronous multi-frame with a linear motion model, denoted by the superscript ℓ. In general, given timestamps $t_1 \leq t_2$, and respective associated poses $\mathbf{T}_1$, $\mathbf{T}_2$, poses at any timestamp t could be linearly interpolated or extrapolated as

$$
\begin{aligned}
\mathbf{T}^{\ell}(t) &= \mathbf{T}_2(\mathbf{T}_2^{-1}\mathbf{T}_1)^{\alpha} \\
&= \mathbf{T}_2\mathrm{Exp}(\alpha\mathrm{Log}(\mathbf{T}_2^{-1}\mathbf{T}_1)), \quad \text{where } \alpha = \frac{t_2 - t}{t_2 - t_1}.
\end{aligned}
$$
(5)

In the context of multi-frames, for each multi-frame MF_i with the representative timestamp $\bar{t}_i$, we define the linear pose parameter $\boldsymbol{\xi}_i^{\ell} \in \mathbb{R}^6$ to represent the minimal 6-DoF robot pose in the world frame at $\bar{t}_i$. It follows that poses at any timestamp t within MF_i could be evaluated with

$$
\begin{aligned}
\mathbf{T}_{wb}^{\ell}(t) &= \mathbf{T}_{wb}^{\ell}(\bar{t}_i)\mathrm{Exp}\left(\frac{\bar{t}_i - t}{\bar{t}_i - \bar{t}_{\mathrm{ref}}}\mathrm{Log}\left(\mathbf{T}_{wb}^{\ell}(\bar{t}_i)^{-1}\mathbf{T}_{wb}^{c}(\bar{t}_{\mathrm{ref}})\right)\right) \\
&= \mathrm{Exp}(\boldsymbol{\xi}_i^{\ell})\mathrm{Exp}\left(\frac{\bar{t}_i - t}{\bar{t}_i - \bar{t}_{\mathrm{ref}}}\mathrm{Log}\left(\mathrm{Exp}(-\boldsymbol{\xi}_i^{\ell})\mathbf{T}_{wb}^{c}(\bar{t}_{\mathrm{ref}})\right)\right),
\end{aligned}
$$
(6)

where $\bar{t}_{\mathrm{ref}}$ and $\mathbf{T}_{wb}^{c}(\bar{t}_{\mathrm{ref}})$ are the respective representative timestamp and evaluated cubic B-spline pose at $\bar{t}_{\mathrm{ref}}$ of the reference multi-frame $\mathrm{MF}_{\mathrm{ref}}$. In practice, new MFs are tracked against a reference *key* multi-frame, so *ref* refers to the MF id of the most recent KMF.

C. Cubic B-Spline Model

We use a cumulative cubic B-spline motion model over key multi-frames to represent the overall trajectory. We use the linear motion model parameters estimated during tracking to initialize cubic B-spline control points, and refine the spline trajectory during mapping and loop closing. In general, given $n + 1$ control points $\boldsymbol{\xi}_0^c, \ldots, \boldsymbol{\xi}_n^c \in \mathbb{R}^6$, and a knot vector $\mathbf{b} \in \mathbb{R}^{n+k+1}$, the cumulative B-spline of order k is defined as:

$$
\mathbf{T}_{wb}^{c}(t) = \mathrm{Exp}\left(\widetilde{B}_{0,k}(t)\boldsymbol{\xi}_0^c\right)\prod_{i=1}^{n}\mathrm{Exp}\left(\widetilde{B}_{i,k}(t)\boldsymbol{\Omega}_i\right),
$$
(7)

where $\boldsymbol{\Omega}_i = \mathrm{Log}\left(\mathrm{Exp}(\boldsymbol{\xi}_{i-1}^c)^{-1}\mathrm{Exp}(\boldsymbol{\xi}_i^c)\right) \in \mathbb{R}^6$ is the relative pose in Lie Algebra twist coordinate form between control poses $\boldsymbol{\xi}_{i-1}^c$ and $\boldsymbol{\xi}_i^c$. The superscript c is used to denote the cubic B-spline motion model. The cumulative basis function $\widetilde{B}_{i,k}(t) = \sum_{j=i}^{n} B_{j,k}(t) \in \mathbb{R}$ is the sum of basis function $B_{j,k}(t)$. Based on the knot vector $\mathbf{b} = \begin{bmatrix} b_0 & \ldots & b_{n+k} \end{bmatrix}$, the basis function $B_{j,k}(t)$ is computed using the de Boor-Cox recursive formula [65], [66], with the base case

$$
B_{p,1}(t) = \begin{cases} 1 & \text{if } t \in [b_p, b_{p+1}] \\ 0 & \text{otherwise} \end{cases}.
$$

For $q \geq 2$,

$$
B_{p,q}(t) = \frac{t - b_p}{b_{p+q-1} - b_p}B_{p,q-1}(t) + \frac{b_{p+q} - t}{b_{p+q} - b_{p+1}}B_{p+1,q-1}(t).
$$

More intuitively, each $\mathbf{T}_{wb}^{c}(t)$ can be interpreted as an n-way interpolation between the control points $\boldsymbol{\xi}_i^c$ with respective interpolation weight $B_{i,k}(t)$. However, instead of directly interpolating $\mathbf{T}_{wb}^{c}(t) = \prod_{i=0}^{n}\mathrm{Exp}(B_{i,k}(t)\boldsymbol{\xi}_i^c)$, we use the cumulative formulation in Eq. 7 for accurate on-manifold interpolation in $\mathbb{SE}(3)$ [45], [64].

Since we use cubic B-splines, $n = 3$ and $k = 4$. In the context of multi-frames, we associate each *key* multi-frame KMF_i with a control point $\boldsymbol{\xi}_i^c$. In addition, since the key multi-frames are not necessarily distributed evenly in time, we cannot utilize a uniform knot vector (as typically employed for modeling rolling-shutter cameras [45] and

LiDARs [54]). Instead, we associate each KMF_i with a *representative timestamp* $\bar{t}_i$ as the median of all image capture times t_{ik} (with the exception of initialization, where $\bar{t}_0$ is defined at the firing time of the overlapping camera pair). We define a non-uniform knot vector according to the representative timestamps. Specifically, we define $\mathbf{b} = \begin{bmatrix} \bar{t}_{i-3} & \bar{t}_{i-2} & \dots & \bar{t}_{i+4} \end{bmatrix} \in \mathbb{R}^8$.

D. Loop Detection

When a new KMF is selected, we run loop detection to check if a previously-seen area is being revisited. For computational efficiency, we only perform loop detection if the most recent successful loop correction took place at least a minimal number of KMFs ago. In our implementation we set this threshold to 30.

For a newly-inserted query KMF_q, the loop candidate KMF_l must pass an odometry check, a multi-camera DBoW3-based [69] similarity check, and a multi-camera geometric verification. We detail this process in the following subsections.

1) Odometry Check: To avoid false loop detection when the robot is staying in the same area, the odometry check ensures that the robot must have traveled a minimal distance since the loop candidate frame. In addition, a minimal time and a minimal number of key frames must have passed since the candidate frame as well. The traveling distance is computed based on the estimated trajectory. The time and key frame count conditions serve as complements in the case when the estimated traveling distance is unreliable. In our experiment, we set the traveling distance threshold to 30m, time to 5s, and the number of KMFs to 30. Note that the KMF threshold in odometry check is different from the KMF threshold described at the beginning of the loop detection section. The former specifies that a candidate KMF must be older than 30 KMFs ago, while the latter dictates that we will only perform loop detection for the current query KMF if the latest loop closing happened at least 30 KMFs ago.

2) Similarity check: For candidates passing the odometry check, we perform a multi-view version of the single-view DBoW3-based similarity check described in ORB-SLAM [4]. The key idea is that images in the loop candidate KMF and the query KMF should have similar appearance. We perform similarity detection with the bag-of-words techniques DBoW3 for place recognition [69]. For each key multi-frame, we concatenate features extracted from all images in the multi-frame to build the DBoW3 vocabulary and compute the DBoW3 vector. Note that the simple concatenation does not take into account of the fact that cameras can be asynchronous, but we argue that the same area should have similar appearance within the short camera firing time interval, and false positives will be filtered by the stricter geometric verification that factors in asynchronous sensors in the next step.

During the similarity check, we first compute the DBoW3 similarity score between the query KMF and the neighboring KMFs that are included in the associated local bundle adjustment window, and we denote the minimum similarity

score as m. Next, we compute the similarity score between the query KMF and all available candidate KMFs and denote the top score as t. Then all the candidates that pass the similarity check must have a similarity score that is greater than $\max(0.01, m, 0.9 * t)$.

3) Geometric check: For each remaining candidate KMF, we perform a geometric check by directly solving for a relative pose between cameras in the candidate KMF and cameras in the query KMF. To identify the camera pairs to be matched, let us consider a setting where the camera rig contains a set of M cameras covering a combined $360°$ FoV, with the cameras denoted as $\mathcal{C}_1, \dots, \mathcal{C}_M$ in the clockwise order. We also assume that the robot is on the ground plane in this setting, *i.e.*, the roll and pitch angles of the robot poses remain the same when the robot revisits the same area. Since a loop can be encountered at an arbitrary yaw angle, there are a total of M possible scenarios of how the multiple cameras between the candidate and the query frame can be matched, where in scenario i, each camera $\mathcal{C}_j$ in the candidate frame is matched to camera $\mathcal{C}_{(j+i)\%M}$ in the query frame.

For each possible matching scenario involving M pairs of cameras, we solve for a discrete relative pose between each camera pair. Specifically, for each pair of cameras, we first perform keypoint-based sparse image matching between the associated image pair by fitting an essential matrix in a RANSAC [85] loop. If the number of inlier matches passes a certain threshold, we associate the inlier matches with the existing 3D map points. Note that different from tracking, here we draw associations in both directions: 2D keypoints in the loop candidate image are associated to 3D map points observed in the query image, and vice versa.

If the number of such keypoint-to-map-point correspondences passes a threshold, we estimate a single relative pose in $\mathbb{SE}(3)$ between the two cameras. Following [4], we perform pose estimation with Horn's method in a RANSAC loop, where within each RANSAC iteration we sample a minimal number of matches, and solve for the discrete pose by minimizing a reprojection error. The hypothesis with the most number of inliers is the final estimate.

The geometric check passes if at least a certain number of camera pairs have a minimum number of inliers for the relative pose estimation. In our full system, we perform geometric check with the $M = 5$ wide cameras covering the surroundings of the vehicle. We consider a geometric check to be successful if there exists a matching scenario where we can successfully estimate the discrete relative pose for at least 2 pairs of cameras, where for each camera pair there are at least 20 inlier correspondence pairs during sparse image matching, 20 pairs of 2D-3D associations, and 20 pairs of inlier correspondences after the relative pose estimation. If there are multiple matching scenarios that pass the check, we select the configuration with the most successfully matched camera pairs and the most inlier correspondences.

The multi-camera geometric verification outputs $\{(\mathcal{C}_{k_l}, \mathcal{C}_{k_q}, \mathbf{T}_{b_{k_q}, b_{k_l}})\}$, which is a set of triplets denoting the camera indices of the matched camera pairs between the loop and the query frames, along with $\mathbf{T}_{b_{k_q}, b_{k_l}}$, which is an

estimated rigid-body transformation from the body frame $\mathcal{F}_b$ at the camera capture time t_{lk_l} to $\mathcal{F}_b$ at time t_{qk_q}.

E. Loop Correction

If a loop candidate KMF_l passes all loop detection checks, we perform loop correction with the geometric verification output. We first build an asynchronous multi-view version of the pose graph in ORB-SLAM [4]. Each node α of the pose graph is encoded by a timestamp t_α representing the underlying robot pose $\mathbf{T}_{wb}(t_\alpha)$. Each edge (α, β) encodes the relative pose constraint $\mathbf{T}_{\beta\alpha}$ representing the rigid transformation from $\mathbf{T}_{wb}(t_\alpha)$ to $\mathbf{T}_{wb}(t_\beta)$.

Specifically, in our pose graph, the nodes are associated with times at $\{\bar{t}_i\}_{\forall KMF_i} \cup \{t_{lk_l}, t_{qk_q}\}_{\forall (k_l, k_q)}$, *i.e.*, the representative timestamps of all existing KMFs, as well as the camera times from matched camera pairs in the geometric verification output. The edges are comprised of: (1) neighboring edges connecting adjacent KMF nodes at times $\bar{t}_{i-1}$ and $\bar{t}_i$, (2) past loop closure edges connecting nodes associated with past query and loop closure KMFs, and (3) the new loop closure edges between nodes at time t_{lk_l} and t_{qk_q}. For edges in case (1) and (2), we compute the relative pose $T_{\beta\alpha}$ by evaluating $(\mathbf{T}_{wb}^c(\bar{t}_\beta))^{-1}\mathbf{T}_{wb}^c(\bar{t}_\alpha)$ based on the current trajectory. For (3), we use the discrete poses $\mathbf{T}_{b_{k_q},b_{k_l}}$ estimated in the geometric verification step in loop detection. Please refer to Fig. 6 for an illustration of the pose graph.

We denote the local KMF windows spanning the query and loop frames as the *welding windows*. In our implementation they are the same size as the local bundle adjustment window. To correct the global drift, we perform a pose graph optimization (PGO) over the continuous-time cubic B-spline trajectory. To better anchor the trajectory, the control poses in the welding window associated to KMF_l are fixed during the pose graph optimization, where the following objective is minimized:

$$E_{PGO}(\{\boldsymbol{\xi}_i^c\}) = E_{rel}(\{\boldsymbol{\xi}_i^c\}) + E_{reg}(\{\boldsymbol{\xi}_i^c\}), \qquad (8)$$

where

$$E_{rel}(\{\boldsymbol{\xi}_i^c\}) = \sum_{(\alpha,\beta)} \rho\left(\left\|\mathbf{e}_{\alpha,\beta}^T(\{\boldsymbol{\xi}_i^c\})\right\|_{\boldsymbol{\Sigma}_{\alpha,\beta}^{-1}}^2\right),$$
$$\text{with } \mathbf{e}_{\alpha,\beta}(\{\boldsymbol{\xi}_i^c\}) = \text{Log}(\mathbf{T}_{\beta\alpha}(\mathbf{T}_{wb}^c(t_\alpha))^{-1}\mathbf{T}_{wb}^c(t_\beta)) \in \mathbb{R}^6 \qquad (9)$$

sums over the relative pose errors of each edge weighted by an uncertainty term $\boldsymbol{\Sigma}_{\alpha,\beta}^{-1}$, and

$$E_{reg}(\{\boldsymbol{\xi}_i^c\}) = \sum_i \rho\left(\left\|\mathbf{r}_i^T(\{\boldsymbol{\xi}_i^c\})\right\|_{\boldsymbol{\Lambda}_i^{-1}}^2\right),$$
$$\text{with } \mathbf{r}_i^T(\{\boldsymbol{\xi}_i^c\}) = \text{Log}\left(\mathbf{T}_i^{-1}\mathbf{T}_{wb}^c(\bar{t}_i)\right) \in \mathbb{R}^6 \qquad (10)$$

is a unary regularization term weighted by uncertainty $\boldsymbol{\Lambda}_i^{-1}$ to anchor each KMF's representative pose at $\mathbf{T}_i = \mathbf{T}_{wb}^c(\bar{t}_i)$ evaluated before the optimization. Empirically, we set the diagonal entries of both $\boldsymbol{\Sigma}_{\alpha,\beta}^{-1}$ and $\boldsymbol{\Lambda}_i^{-1}$ to 1.0. The regularization term helps to better constrain the optimization especially when there is a large loop (*i.e.*, q is much bigger than l) and a large amount of drift to correct. ρ is the robust

norm and we again use the Huber loss in practice. The energy term is minimized with the LM algorithm.

After PGO, we next update the map points with the adjusted trajectory. If a map point is observed in multiple images, we update the map point position with the median of all pose corrections related to the map point. Note that ideally, we would want to solve a global bundle adjustment problem that jointly refines the entire trajectory and all map points at the same time. However, with a long trajectory and many observations from multi-view cameras, global bundle adjustment becomes computational expensive or even infeasible. The two-stage process described above, where we first optimize the poses and then update the map points, is a light-weight approximation that is sufficient under most circumstances.

Furthermore, note some new map points in the query KMF window have been created during recent local mappings, but they may correspond to points already triangulated in the previous pass through the revisited area. As a result, we next deduplicate the re-triangulated map points. We first match image pairs in the candidate KMF welding window and the query KMF welding window to identify and fuse these map points. We then perform a local bundle adjustment over the motion and map points corresponding to the two welding windows. The purpose of the local bundle adjustment is to refine both map points and control poses in the query KMF window. To anchor candidate poses, we freeze the control poses corresponding to the candidate welding window in the optimization.

C. DATASET

A. Existing SLAM Benchmarks

As described in the main paper, existing SLAM datasets fall short in terms of geographic diversity, modern sensor layouts, or scale. In this section, we describe the most relevant modern SLAM datasets together with their primary limitations.

The KITTI Odometry Benchmark [6] covers 40km of driving through Karlsruhe, Germany using a vehicle equipped with a stereo camera pair, a 64-beam LiDAR, IMU, and RTK-based ground truth. However, most sequences in the dataset have relatively small numbers of dynamic objects, and all the data is captured in sunny or overcast weather, which is not representative of the variety of conditions which can be encountered by commercial AVs. The NCLT dataset [74] covers a larger distance in the University of Michigan campus using a custom-built robotic Segway equipped with three LiDARs an IMU and an omnidirectional camera. While the scale of the dataset is large, its geographic diversity is lacking, being constrained to a university campus, while at the same time not capturing the same challenging motion patterns which would be encountered by an SDV.

The Oxford RobotCar [70] dataset covers over 1000km of driving in challenging conditions containing a large number of dynamic objects as well as strong weather and lighting variation. However, since it is focused on a single primary trajectory it lacks geographic diversity. Moreover, it does not provide 360° camera data in HD, which is critical for
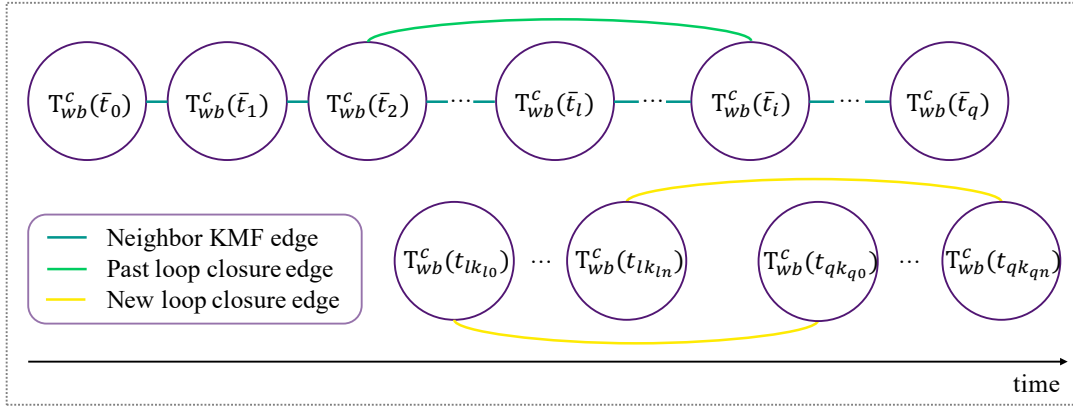
9

Fig. 6: Illustration of the loop correction pose graph. The nodes correspond to robot poses at the representative timestamps of all KMFs + capture times of the matched cameras in the new loop closure KMFs. The edges consist of neighboring KMF edges, past loop closing edges, and the new multi-view loop closing edges. If n camera pairs are successfully matched during the loop detection stage, then there should be n new loop closing edges.
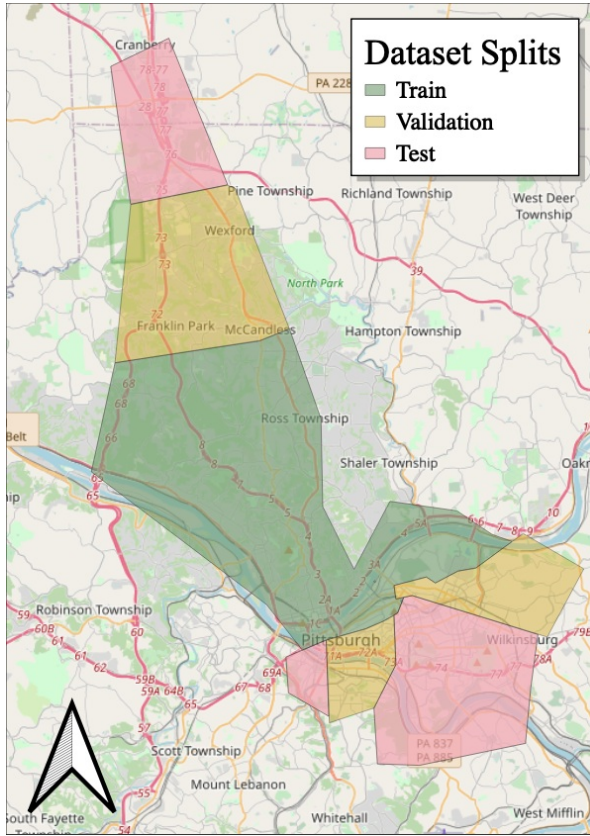

Fig. 7: The geographic splits of the proposed datasets. They are designed such that train, val, and test each covers a balanced blend of environment types (highway, urban, industrial, residential, etc.). The splits are also selected to have similar distributions of weather, loop closures, etc.

SDVs. Similarly, the Ford Multi-AV Dataset [67] contains a large volume of data but is focused on a relatively limited 60km route which is traversed repeatedly by multiple SDVs. The A2D2 Dataset [68] contains a high-resolution multi-camera multi-LiDAR setup optimized for 3D reconstruction. However, the approximate total scale of A2D2 is still on the order of a few hours of driving, which is insufficient for evaluating robust SLAM system in a wide variety of challenging conditions.

Finally, the recent 4Seasons dataset [32] covers diverse areas in a wide range of environments (highway, industrial, residential, etc.) over a long time period, but lacks the HD multi-view sensor array common in commercial SDVs.

Note that in the dataset overview table in the main paper we label A2D2 and Ford Multi-AV as non-asynchronous dataset on the basis that their cameras are not described as following the LiDAR or any custom firing pattern causing more than 2–3ms of delay. While Ford Multi-AV does have cameras firing at different frame rates, with the higher-resolution front stereo operating at 5Hz and the other cameras at 15Hz, we do not consider this as a true asynchronous setting.

Additionally, even though datasets such as Waymo Open [18] and nuScenes [19] include asynchronous cameras, they are focused on perception tasks and contain short sequences (e.g., less than a minute each). Therefore, we do not consider them in our evaluation as they are too short to robustly evaluate SLAM algorithms.

B. Dataset Details

The dataset has been selected using a semi-automatic curation process to ensure all splits cover similar categories of environments, while at the same time having substantial appearance differences between the splits.

Table IV shows the high-level statistics of the train, validation, and test partitions of the dataset. Figure 7 shows the geographic regions of the splits.

The cameras are all RGB and have a resolution of 1920×1200 pixels, a global shutter; the (asynchronous) shutter

TABLE IV: The proposed dataset, AMV-Bench, in numbers.

Split	Sequences	Distance (km)	Time (h)
Train	65	281	12
Validation	25	103	4
Test	26	98	5
Total	116	482	21

timestamps are recorded with the rest of the data. The wide-angle and narrow-angle cameras have rectified FoVs of $76.6° \times 52.6°$ and $37.8° \times 24.2°$, respectively.

Furthermore, the 6-DoF ground-truth was generated using an offline HD-map based localization system, which enables SLAM systems to be evaluated at the centimeter level.

D. ADDITIONAL EXPERIMENTS

A. Full Implementation Details

All images are downsampled to 960×600 for both our method as well as all baselines. In our system, we extract 1000 ORB [76] keypoints from each image, using grid-based sampling [4] to encourage homogeneous distribution. Image matching is performed with nearest neighbor + Lowe's ratio test [77] with a ratio threshold of 0.7. The initial 2D matches between each image pair are additionally filtered by fitting an essential matrix with RANSAC. The inlier correspondences are used (1) as input to the multi-view PnP during tracking, (2) for new map point triangulation during mapping, (3) for geometric verification during loop closing.

In our system, we use the synchronous stereo camera pair during initialization. During tracking, we match image pairs captured by the same camera. During new map point creation, we match images captured by the same camera between the new KMF and four previous KMFs, and triangulate new map points based on the 2D matches. We additionally triangulate new map points from the stereo cameras within the new key multi-frame. Note that we do not match between different wide cameras in the same multi-frame due to little overlap between them and ORB's reduced performance in wide baseline image matching settings. Please see Sec. D-D.1 for details.

During tracking, we randomly sample 7 pairs of correspondences within each PnP RANSAC loop to solve for a hypothesis. The pose estimate hypothesis with the most number of inliers becomes the final estimate.

During key frame selection, a new KMF is inserted either (1) when the estimated local translation against reference KMF is over 1m, or local rotation is over $1°$, or (2) when under 35% of the map points are re-observed in at least two camera frames, or (3) when a KMF hasn't been inserted for 20 consecutive MFs. Note that (3) is necessary to model the spline trajectory when the robot stays stationary.

We perform bundle adjustment over a recent window of size $N = 11$. We cull map points with reprojection error over 1.5 pixels.

Following [4], the uncertainty weighting Σ in both tracking and bundle adjustment is based on the scale level where ORB features are extracted. Keypoints extracted from larger/coarser scale levels are less precise and therefore correspond to higher uncertainty and lower weighting during pose estimation.

During the pose graph optimization in loop closure, the uncertainty weighting for the relative pose error terms and the regularization terms are all set to 1.0.

In our system and all our ablation implementations, we declare a tracking failure if the estimated pose parameters yield under 12 inlier PnP correspondences in total. We declare bundle adjustment failure and not apply the bundle adjustment update if after bundle adjustment any of the pose parameters is changed by more than 6 meters or 20 degrees. We selected these values empirically based on training set performance. We stop the system in the middle of processing a sequence if there are at least 5 successive tracking failures, or if there are at least 5 successive bundle adjustment failures.

We use the Ceres Solver [86] for modeling and solving the non-linear optimization problems arising in tracking, bundle adjustment, and loop closure.

B. Third-Party Baseline Experiment Details

For all ORB-SLAM [4], [5] experiments, we lowered the default tracking failure threshold from 30 matching inliers to 10 matching inliers. This is to increase the tolerance for tracking failures, as the system with 10–30 matching inliers was able to complete larger portions of the training sequences without much compromise of local tracking errors. Apart from the inlier threshold,, we use all other default hyperparameters provided for the KITTI experiments by the authors[a], which extract 2000 ORB keypoints per image, while we only extract 1000 ORB keypoints for our system and all our baseline and ablation study implementations.

We use the default KITTI hyperparameters for LDSO [75] provided by the authors[b].

For the keypoint extractor ablation study, we extract 1000 keypoints and associated features from each image for all methods. We match features with Lowe's ratio test. The ratio is tuned on the training set. We use 0.7 for ORB [76] and RootSIFT [77], [84], and 0.8 for SuperPoint [79] and R2D2 [83]. For SuperPoint, we run the provided pre-trained model [c]. For R2D2, we run the provided pre-trained `r2d2_WASF_N16` model [d].

C. Metrics

In the additional experiments, aside from reporting the median and AUC for the aggregated ATE and RPE results, we also report the ATE and RPE errors at the 90th percentile, *i.e.* x with $f(x) = 0.9$ where f is the cumulative error curve. Compared to median (the 50th percentile), the 90th percentile metric better characterizes outlier behavior, and the AUC metric gives a better characterization of the overall performance, while being able to model system failures.

[a]https://github.com/raulmur/ORB_SLAM2
[b]https://github.com/tum-vision/LDSO
[c]https://github.com/rpautrat/SuperPoint
[d]https://github.com/naver/r2d2

11

D. Quantitative Results

In this subsection we show additional details on the main paper results, as well as additional ablation studies.

1) Detailed main paper results: Table V and Table VI compare (1) third-party baselines, (2) our implementation of the synchronous baselines, and (3) our main system, in the SLAM mode and visual odometry (VO) mode respectively. In the VO mode we disable loop closing, and relocalization in ORB-SLAM. Figs. 8 and 9 plot the respective cumulative error curves.

Note that some metrics of our full system with loop closure in Table V are slightly worse than those without loop closure in Table VI. The difference can also be observed in the DSO experiments. The difference in our system is due to the stochasticity of the trials, instead of loop closing failures. To support our claim, in the main paper we plot the drift relative to the ground truth with and without loop closure to show that loop closing successfully reduced global drifts at all key multi-frames where loop closing was performed. Furthermore, Table VII compares all metrics on the 8/25 validation sequences where loop closure was performed, and shows that the 8 loop closing sequences did not contribute to the metrics differences over all 25 validation sequences.

Table VIII showcases the motion model ablation study results. For simplicity, loop closure is disabled. The asynchronous linear motion model setting represents the trajectory with a linear motion model parameterized by a 6-DoF pose $\boldsymbol{\xi}_i^\ell$ at each representative timestamp $\bar{t}_i$. The motion model is explained in detail in Sec. B. Similar to the main system, we estimate linear motion model parameter during tracking, and we jointly refine the linear motion model parameters along with the map points during bundle adjustment, with the reprojection energy similar to that of tracking. Our experiments show that modeling the vehicle motion using cubic B-splines leads to improved performance due to the splines' ability to impose a realistic motion prior on the estimated trajectories.

In the additional motion model ablation results, we also compare with linear and cubic B-spline models that perform interpolation in $\mathbb{SO}(3)$ and $\mathbb{R}^3$ separately, instead of jointly in $\mathbb{SE}(3)$. Previous works [46], [59], [87], [88] have shown that the split interpolation formulation is generally better in terms of both computation time and trajectory representation.

The results show our main system with cubic B-spline model and full interpolation in $\mathbb{SE}(3)$ has the best performance overall. The split-interpolation motion models have slightly worse performance in our experiments, most likely because full interpolation in $\mathbb{SE}(3)$ is more apt at modeling curvy trajectories (*e.g.*, during turns).

Table IX and Fig. 11 compare with additional camera configurations during tracking and mapping. In the additional camera configurations, during local mapping we additionally match and triangulate new map points between the wide front left and wide front middle cameras, and between the wide front middle and wide front right cameras within the same key multi-frame. Note that in the settings without stereo cameras, we still use stereo cameras (only) for initialization.

Configurations without stereo cameras have worse performance, and we argue this is due to ORB's poor performance in the much harder wide-baseline image matching problem posed by little overlap between the wide front cameras during new map point creation. The table shows that with keypoint extractors such as SuperPoint [79] in place of ORB, this performance gap is significantly narrowed.

Table X and Fig. 12 show the full results of the keypoint extractor ablation study. Compared to ORB, SIFT and SuperPoint finish more sequences and have better ATE, with the caveat that feature extraction takes more time.

2) Per sequence results: Table XI shows per-sequence errors comparing baselines and our method, all with loop closure. We report the mean over all three trials. If at least one trial did not complete the sequence successfully, we do not report results for that sequence.

3) KMF heuristics ablation: To study the effect of the combined KMF selection heuristics that factors in both map point reobservability and motion, we perform an ablation study where we run our stereo + synchronous discrete-time motion model implementation with a reobservability-only KMF heuristic. Table XII shows that our stereo implementation with a reobservability-only heuristic performs worse than the stereo system with the more robust combined heuristic. Table XIII compares the number of key frames inserted by ORB-SLAM2, our stereo sync with a reobservability-only heuristic, and our stereo sync with a reobservability+motion heuristic. The key frame numbers are taken from a randomly-selected trial. The table shows that our reobservability-only heuristic in general inserts fewer key frames than ORB-SLAM2, and that the combined heuristic selects almost twice as many key frames during highway sequences, when the vehicle is driving very fast in a highly repetitive scene.

E. Qualitative Results

1) Qualitative Trajectories:

a) Full Results for Ours-A vs. ORB-SLAM2: Fig. 13 and Fig. 15 showcase the trajectories in all 25 validation sequences, comparing ORB-SLAM2 using only the stereo cameras to our full system using all 7 asynchronous cameras. Fig. 14 depicts the trajectories ins all 65 training sequences.

In the following paragraphs, we qualitatively showcase the failure cases of our main paper ablation study baselines.

b) Motion Model Ablation: Fig. 16 plots failure cases of the linear motion model and the discrete-time motion model with a wrong synchronous assumption. The linear motion model trial failed early due to repeated mapping failures in a challenging case with dynamic objects, and the synchronous model had huge estimation errors during complex maneuvers such as reversing and parking.

c) Camera Ablation: Fig. 17 plots trajectories estimated with different camera configurations, highlighting failure cases resulted from camera configurations with a narrower field of view in challenging conditions like view obstruction, low light, rainy environments and low-textured highway driving.

TABLE V: Baseline methods. M=monocular, S=stereo, A=all cameras.

Method	RPE-T (cm/m)			RPE-R (rad/m)			ATE (m)			SR (%)
	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	
LDSO-M [75]	42.72	-	28.08	8.02E-05	-	54.23	594.39	-	44.67	62.67
ORB-M [4]	34.00	-	25.66	5.49E-05	-	63.77	694.37	-	42.65	64.00
ORB-S [5]	1.85	-	65.70	3.29E-05	-	70.47	30.74	-	74.31	77.33
Sync-S	<u>1.30</u>	-	<u>77.54</u>	<u>2.91E-05</u>	-	<u>78.37</u>	<u>24.53</u>	-	<u>77.44</u>	<u>84.00</u>
Sync-A	2.15	-	68.46	3.47E-05	-	70.47	58.18	-	75.01	74.67
Ours-A	**0.35**	**1.99**	**88.63**	**1.13E-05**	**6.50E-05**	**88.17**	**6.13**	**322.95**	**88.82**	**92.00**
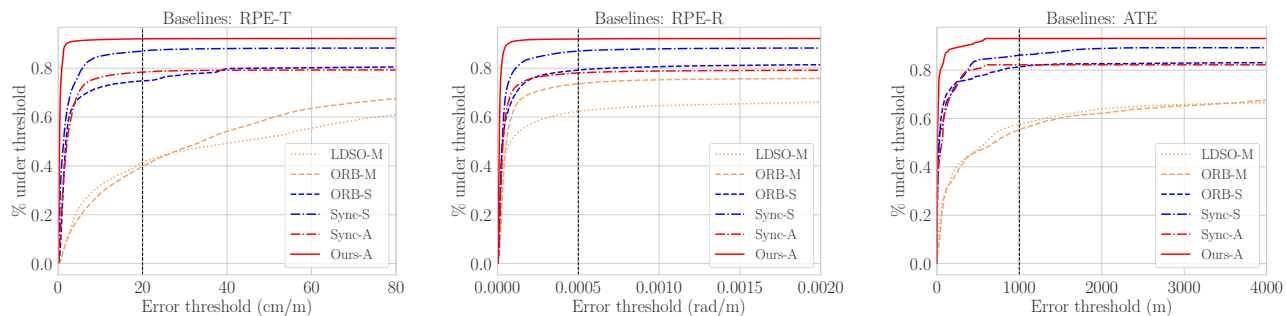
Fig. 8: Cumulative error curves comparing all baseline methods and our full system, with loop closure.

TABLE VI: Baseline methods, all in the visual odometry (VO) mode with loop closing (and relocalization in ORB-SLAM) disabled.

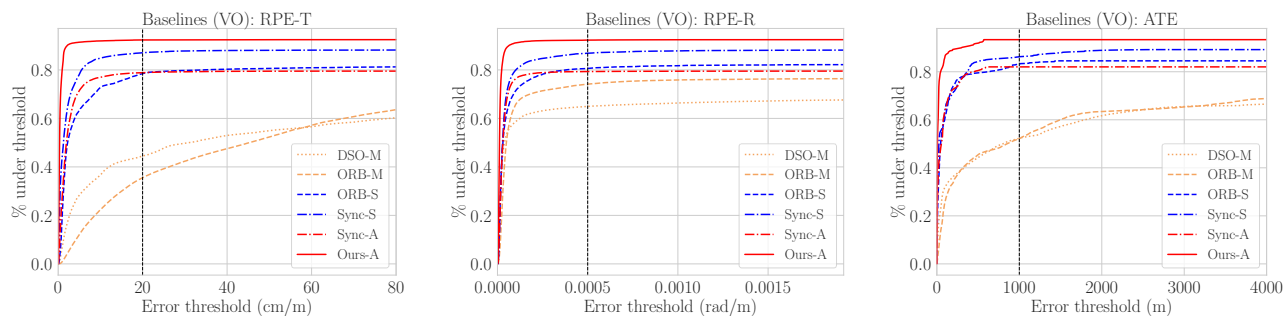| Method	RPE-T (cm/m)			RPE-R (rad/m)			ATE (m)			SR (%)
	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	
DSO-M [3]	30.99	-	32.93	3.88E-05	-	58.98	801.99	-	41.87	64.00
ORB-M (VO) [4]	45.22	-	20.33	4.93E-05	-	64.91	849.64	-	40.62	58.67
ORB-S (VO) [5]	2.24	-	64.91	2.99E-05	-	72.30	45.28	-	74.61	66.67
Sync-S (VO)	<u>1.27</u>	-	<u>77.77</u>	<u>2.80E-05</u>	-	<u>78.72</u>	<u>24.07</u>	-	<u>77.64</u>	<u>85.33</u>
Sync-A (VO)	1.97	-	69.46	2.96E-05	-	73.39	55.24	-	75.11	70.67
Ours-A (VO)	**0.35**	**2.14**	**88.79**	**1.11E-05**	**6.30E-05**	**88.47**	**6.53**	**299.30**	**89.04**	**92.00**

Fig. 9: Cumulative error curves comparing all baseline methods and our full system, without loop closure.

TABLE VII: Ablation study on loop closure on 8 validation sequences where loop closing was performed.

Method	RPE-T (cm/m)			RPE-R (rad/m)			ATE (m)			SR (%)
	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	
Ours-A (VO)	**0.27**	0.99	96.67	**9.51E-06**	**3.07E-05**	95.73	3.87	25.08	97.67	**100.00**
Ours-A	0.28	**0.92**	**96.78**	9.83E-06	3.19E-05	**95.88**	2.97	20.58	97.73	**100.00**

13

TABLE VIII: Ablation study on motion models. All experiments in visual odometry (VO) mode with loop closing disabled. Split indicates interpolation in $\mathbb{SO}(3)$ and $\mathbb{R}^3$ separately [59] instead of jointly in $\mathbb{SE}(3)$.

Method	RPE-T (cm/m)			RPE-R (rad/m)			ATE (m)			SR (%)
	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	
Sync Assumption	1.97	-	69.46	2.96E-05	-	73.39	55.24	-	75.11	70.67
Linear (Split)	0.36	2.51	88.08	**1.10E-05**	7.96E-05	87.79	6.25	580.23	87.43	**92.00**
Linear	0.41	2.71	87.76	1.11E-05	**5.99E-05**	88.39	6.09	429.86	88.31	89.33
Cubic B-Spline (Split)	0.38	3.34	87.86	1.12E-05	9.01E-05	87.69	**5.15**	588.96	87.34	**92.00**
Cubic B-Spline	**0.35**	**2.14**	**88.79**	1.11E-05	6.30E-05	**88.47**	6.53	**299.30**	89.04	**92.00**
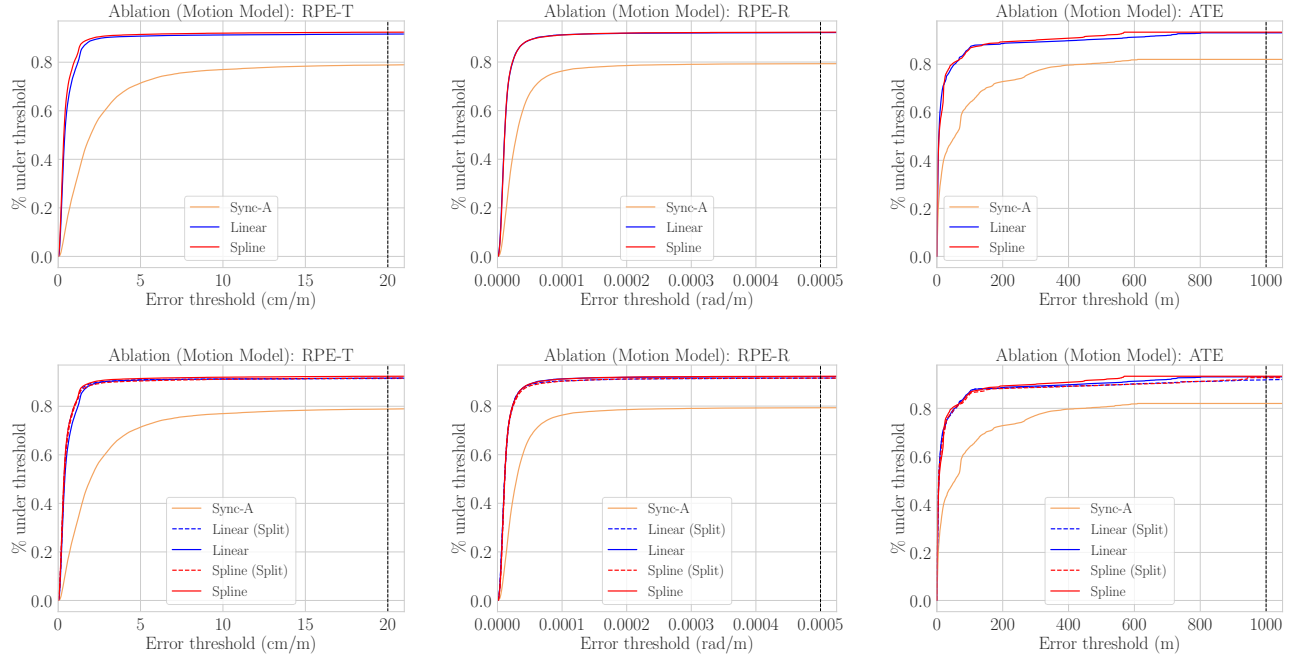
Fig. 10: Cumulative error curves of the motion model ablation study. (Top) The three comparisons in the main paper. (Bottom) All comparisons in the additional experiments.

TABLE IX: Ablation study on camera rigs in the VO mode, all initialized with the stereo cameras. s = stereo, wf = wide-front, wb = wide-back, ✓ is used for intra-frame new map point creation during mapping. The last row in the ORB table represents the main system.

|	Camera Config			RPE-T (cm/m)			RPE-R (rad/m)			ATE (m)			SR (%)
	s	wf	wb	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	
ORB [76]	✓̌			0.70	-	79.86	1.93E-05	-	80.48	11.44	-	75.75	88.00
	✓̌	✓		0.41	8.52	84.88	1.21E-05	2.48E-04	85.86	9.00	802.88	84.92	90.67
		✓̌		6.07	-	49.00	4.58E-05	-	52.96	53.76	-	55.05	57.33
		✓̌	✓	1.16	-	63.94	1.65E-05	-	74.34	18.63	-	72.10	74.67
	✓̌	✓̌	✓	0.36	3.43	88.43	1.12E-05	**5.35E-05**	88.60	5.95	298.08	89.05	**92.00**
	✓̌	✓	✓	**0.35**	2.14	**88.79**	**1.11E-05**	6.30E-05	88.47	6.53	299.30	89.04	**92.00**
SuperPoint [79]	✓̌			0.64	19.44	82.97	1.59E-05	5.77E-04	83.09	16.40	463.41	85.09	82.67
	✓̌	✓		0.44	1.62	92.83	**1.01E-05**	3.48E-05	93.43	8.33	106.31	93.92	97.33
		✓̌		1.06	-	78.32	2.88E-05	-	76.73	21.08	-	73.48	88.00
		✓̌	✓	0.54	1.85	89.42	1.84E-05	6.51E-05	88.31	10.62	413.52	87.67	96.00
	✓̌	✓̌	✓	**0.38**	**1.22**	**95.66**	1.04E-05	2.63E-05	95.86	**5.38**	86.67	**96.38**	**100.00**
	✓̌	✓	✓	0.41	1.28	95.28	1.03E-05	**2.54E-05**	95.90	6.83	**78.06**	96.34	98.67
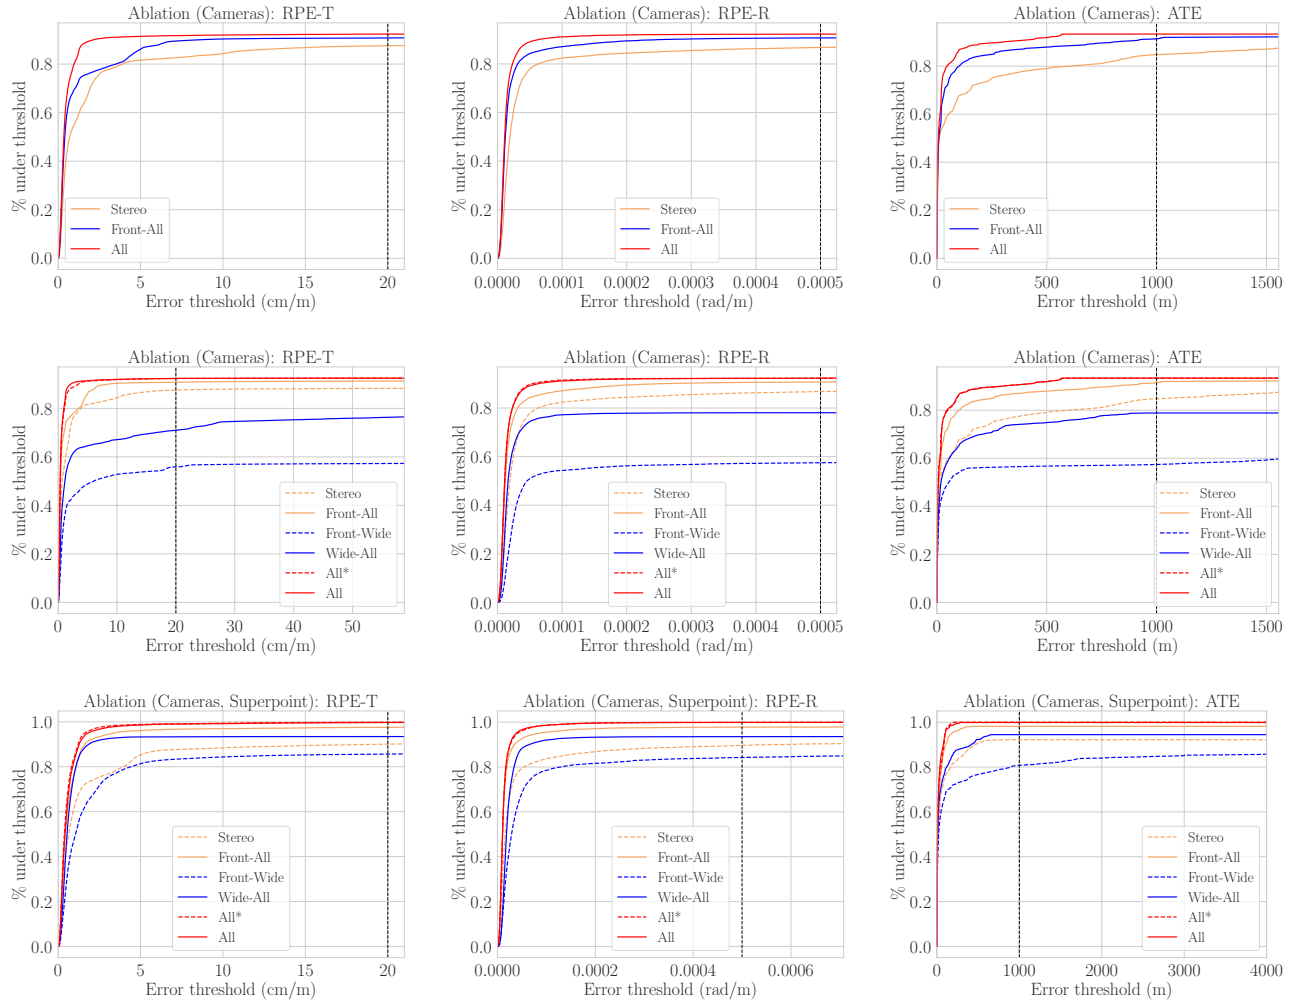
Fig. 11: Cumulative error curves of the camera ablation study. (Top) The three configurations in the main paper. (Middle) All camera configurations in the additional experiments. (Bottom) All camera configurations with SuperPoint in place of ORB as the keypoint extractor in the additional experiments. Camera configuration legend order corresponds to the order in the table.

TABLE X: Ablation study for keypoint extractors in the VO mode. Time is the average feature extraction time per image using 24 CPU cores.

| Method | Time (s) | RPE-T (cm/m) | | | RPE-R (rad/m) | | | ATE (m) | | | SR (%) |
		@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	
RootSIFT [84]	0.10	0.41	2.26	94.06	1.11E-05	2.69E-05	95.84	6.66	123.72	91.50	**98.67**
SuperPoint [79]	0.35	0.41	**1.28**	**95.28**	**1.03E-05**	**2.54E-05**	95.90	6.83	**78.06**	**96.34**	**98.67**
R2D2 [83]	20.50	0.41	-	84.09	1.40E-05	-	83.72	7.42	-	86.62	88.00
ORB [76]	**0.01**	**0.35**	2.14	88.79	1.11E-05	6.30E-05	88.47	**6.53**	299.30	89.04	92.00

d) Keypoint Extractor Ablation: Fig. 18 plots trajectories estimated by SLAM systems that use ORB, RootSIFT [84] and SuperPoint [79] respectively as the keypoint extractor. RootSIFT and SuperPoint trajectories visually align better with the ground truth and are able to complete more challenging rainy highway sequences.

2) Loop Closure: Figure 19 shows a failure case consisting in a false positive loop detection in the stereo setting. The large bus dominates the field of view of both cameras while also having rich texture due to the lights, ad, *etc.*, causing a loop to be incorrectly closed. Multi-view loop closure correctly rejects this case and many similar others. This highlights the importance of multiple cameras for robust SLAM in the real world. For more qualitative results on the loop closure in the main system, please refer to the supplementary video.

3) Qualitative Map: Figures 20, 21, 22, and 23 showcase visualizations of some of the maps produced by AMV-SLAM. Please refer to our supplementary video for additional qualitative results.

TABLE XI: Per sequence errors of all baselines and our method. Errors averaged over all three trials at all evaluated timestamps. - denotes at least one trial did not successfully complete the sequence. RPE-T(cm/m), RPE-R(rad/m), ATE(m).

| | Monocular | | | | | | Stereo | | | | | | All-Camera | | | | | |
| | LDSO-M | | | ORB-M | | | ORB-S | | | Sync-S | | | Sync-A | | | Ours-A | | |
Sequence	RPE-T	RPE-R	ATE	RPE-T	RPE-R	ATE	RPE-T	RPE-R	ATE	RPE-T	RPE-R	ATE	RPE-T	RPE-R	ATE	RPE-T	RPE-R	ATE
day_no_rain_0	-	-	-	-	-	-	2.95	8.44E-05	40.96	0.51	1.70E-05	6.17	-	-	-	**0.46**	**1.25E-05**	7.38
day_no_rain_1	46.45	1.61E-03	225.75	1.25	2.78E-05	5.43	1.29	3.07E-05	5.09	0.90	6.57E-05	2.74	1.66	6.55E-05	8.43	**0.22**	**1.00E-05**	**0.55**
day_no_rain_2	-	-	-	5.46	4.89E-05	52.75	1.72	7.37E-05	12.04	0.41	2.45E-05	5.54	-	-	-	**0.29**	**1.20E-05**	**3.53**
day_no_rain_3	21.63	2.63E-05	382.04	-	-	-	1.53	3.33E-05	13.03	0.33	2.00E-05	2.95	0.85	1.71E-05	5.18	**0.17**	**1.16E-05**	**2.71**
day_no_rain_4	-	-	-	31.93	4.33E-05	426.90	**0.64**	**2.17E-05**	8.85	0.96	3.61E-05	10.17	-	-	-	0.96	2.57E-05	**3.74**
day_no_rain_5	-	-	-	-	-	-	1.06	2.77E-05	6.64	1.25	5.11E-05	6.49	2.61	5.28E-05	7.83	**0.44**	**1.55E-05**	**1.18**
day_no_rain_6	4.11	1.13E-04	**17.04**	1.74	4.91E-05	47.99	-	-	-	0.79	4.01E-05	24.63	3.27	2.00E-04	73.84	**0.24**	**1.24E-05**	20.11
day_no_rain_7	0.72	3.35E-05	1.75	2.74	1.52E-04	50.96	0.64	**1.58E-05**	1.41	0.28	1.84E-05	**0.38**	0.50	1.61E-05	1.80	0.23	1.64E-05	0.49
day_no_rain_8	-	-	-	44.84	7.88E-05	411.75	-	-	-	1.82	4.80E-05	26.81	1.63	3.13E-05	23.36	**0.40**	**1.28E-05**	**5.30**
day_no_rain_9	17.47	4.31E-04	74.46	26.15	1.47E-04	571.95	-	-	-	0.37	1.82E-05	2.70	-	-	-	**0.33**	**1.70E-05**	**2.07**
day_no_rain_10	31.34	4.76E-05	110.64	14.70	3.54E-05	73.36	1.12	3.73E-05	7.43	0.33	1.66E-05	2.91	0.41	1.55E-05	2.66	**0.27**	**9.05E-06**	**1.96**
day_no_rain_11	1.18	2.50E-05	11.79	2.47	3.16E-05	82.23	0.73	2.50E-05	3.12	0.26	1.53E-05	**1.70**	0.28	1.49E-05	2.86	**0.24**	**1.12E-05**	1.38
day_rain_0	23.37	2.41E-05	652.45	-	-	-	-	-	-	1.59	5.73E-05	65.56	4.06	2.88E-05	73.63	**0.31**	**1.33E-05**	**12.56**
day_rain_1	-	-	-	19.60	3.43E-05	228.22	-	-	-	0.49	3.33E-05	7.85	-	-	-	**0.31**	**1.52E-05**	**3.01**
day_rain_2	8.45	2.75E-05	34.99	19.52	8.92E-05	82.56	11.85	6.23E-04	130.85	-	-	-	-	-	-	**0.77**	**2.05E-05**	**4.91**
day_rain_3	-	-	-	31.31	1.11E-04	219.47	-	-	-	-	-	-	-	-	-	**0.37**	**1.15E-05**	**14.46**
day_rain_4	23.91	3.19E-04	138.17	29.49	1.94E-04	173.28	5.19	1.79E-04	27.15	6.59	4.39E-05	83.15	-	-	-	**0.72**	**2.14E-05**	**7.21**
day_rain_5	11.88	2.40E-05	25.51	4.65	2.71E-05	8.62	10.32	1.92E-04	18.70	0.44	3.60E-05	0.86	1.86	3.28E-05	2.90	**0.16**	**1.23E-05**	**0.51**
hwy_no_rain_0	-	-	-	-	-	-	-	-	-	2.21	4.02E-05	218.65	2.51	3.25E-05	190.31	**0.41**	**1.52E-05**	**52.94**
hwy_no_rain_1	-	-	-	-	-	-	2.92	3.26E-05	49.79	3.21	9.19E-05	888.29	2.91	4.01E-05	92.70	**0.59**	**1.49E-05**	**22.54**
hwy_no_rain_2	-	-	-	-	-	-	-	-	-	3.08	3.30E-05	324.44	3.34	3.56E-05	395.12	**0.45**	**1.40E-05**	377.00
hwy_no_rain_3	-	-	-	-	-	-	9.04	7.69E-05	852.20	5.53	7.38E-05	185.56	3.38	5.60E-05	84.41	**0.48**	**1.50E-05**	**36.58**
hwy_rain_0	-	-	-	-	-	-	-	-	-	0.74	2.57E-05	171.41	1.38	1.86E-05	168.77	**0.31**	**9.69E-06**	**63.98**
hwy_rain_1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
hwy_rain_2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

TABLE XII: Ablation study for the impact of the KMF-selection heuristics on the system performance.

| Method | KMF Heuristics | | RPE-T (cm/m) | | | RPE-R (rad/m) | | | ATE (m) | | | SR (%) |
	reobservability	motion	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	@0.5	@0.9	AUC(%)	
Sync-S	✓		2.90	-	65.24	3.28E-05	-	71.39	68.56	-	68.02	82.67
Sync-S	✓	✓	**1.30**	-	**77.54**	**2.91E-05**	-	**78.37**	**24.53**	-	**77.44**	**84.00**

TABLE XIII: Number of KMFs selected per validation sequence, comparing ORB-SLAM2, ours stereo with reobservability-only heuristics, and ours stereo with a combined heuristics. Empty cells correspond to unfinished sequences.

| sequence	ORB-S [5]	Ours-S (r-only)	Ours-S (combined)
day_no_rain_0	1718	1322	1759
day_no_rain_1	2249	1645	2054
day_no_rain_2	3795	2370	2690
day_no_rain_3	1245	969	1399
day_no_rain_4	2364	1734	2143
day_no_rain_5	-	1455	2034
day_no_rain_6	3373	-	2983
day_no_rain_7	501	464	633
day_no_rain_8	1724	1121	2104
day_no_rain_9	-	1598	1744
day_no_rain_10	1263	982	1505
day_no_rain_11	1593	1107	1559
day_rain_0	-	1522	2838
day_rain_1	-	1902	3043
day_rain_2	1097	-	-
day_rain_3	1952	2181	-
day_rain_4	2520	2339	3119
day_rain_5	383	408	531
hwy_no_rain_0	3646	3507	6532
hwy_no_rain_1	2400	2094	3994
hwy_no_rain_2	3358	2495	4632
hwy_no_rain_3	2124	2242	4484
hwy_rain_0	1835	2316	5132
hwy_rain_1	-	-	-
hwy_rain_2	-	-	-
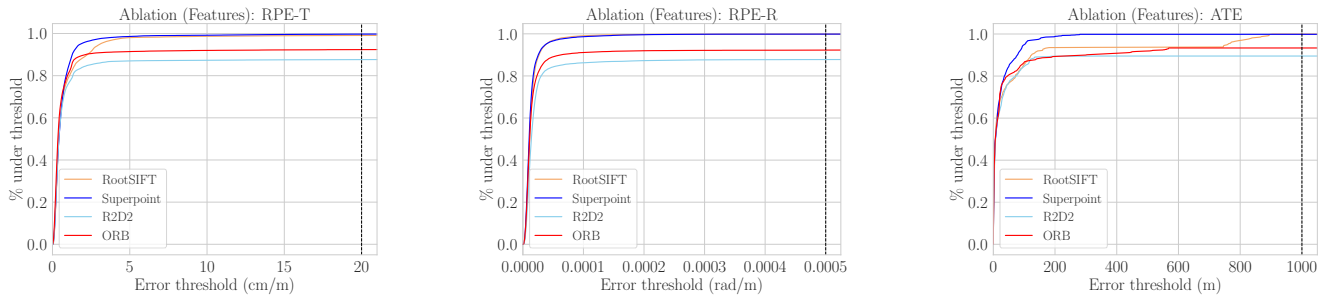
Fig. 12: Cumulative error curves of the keypoint extractor ablation study. Our experiments show that while ORB features still remain competitive, SuperPoint features lead to the best overall performance, especially in terms of translational error. However, this comes at a much higher computational cost.
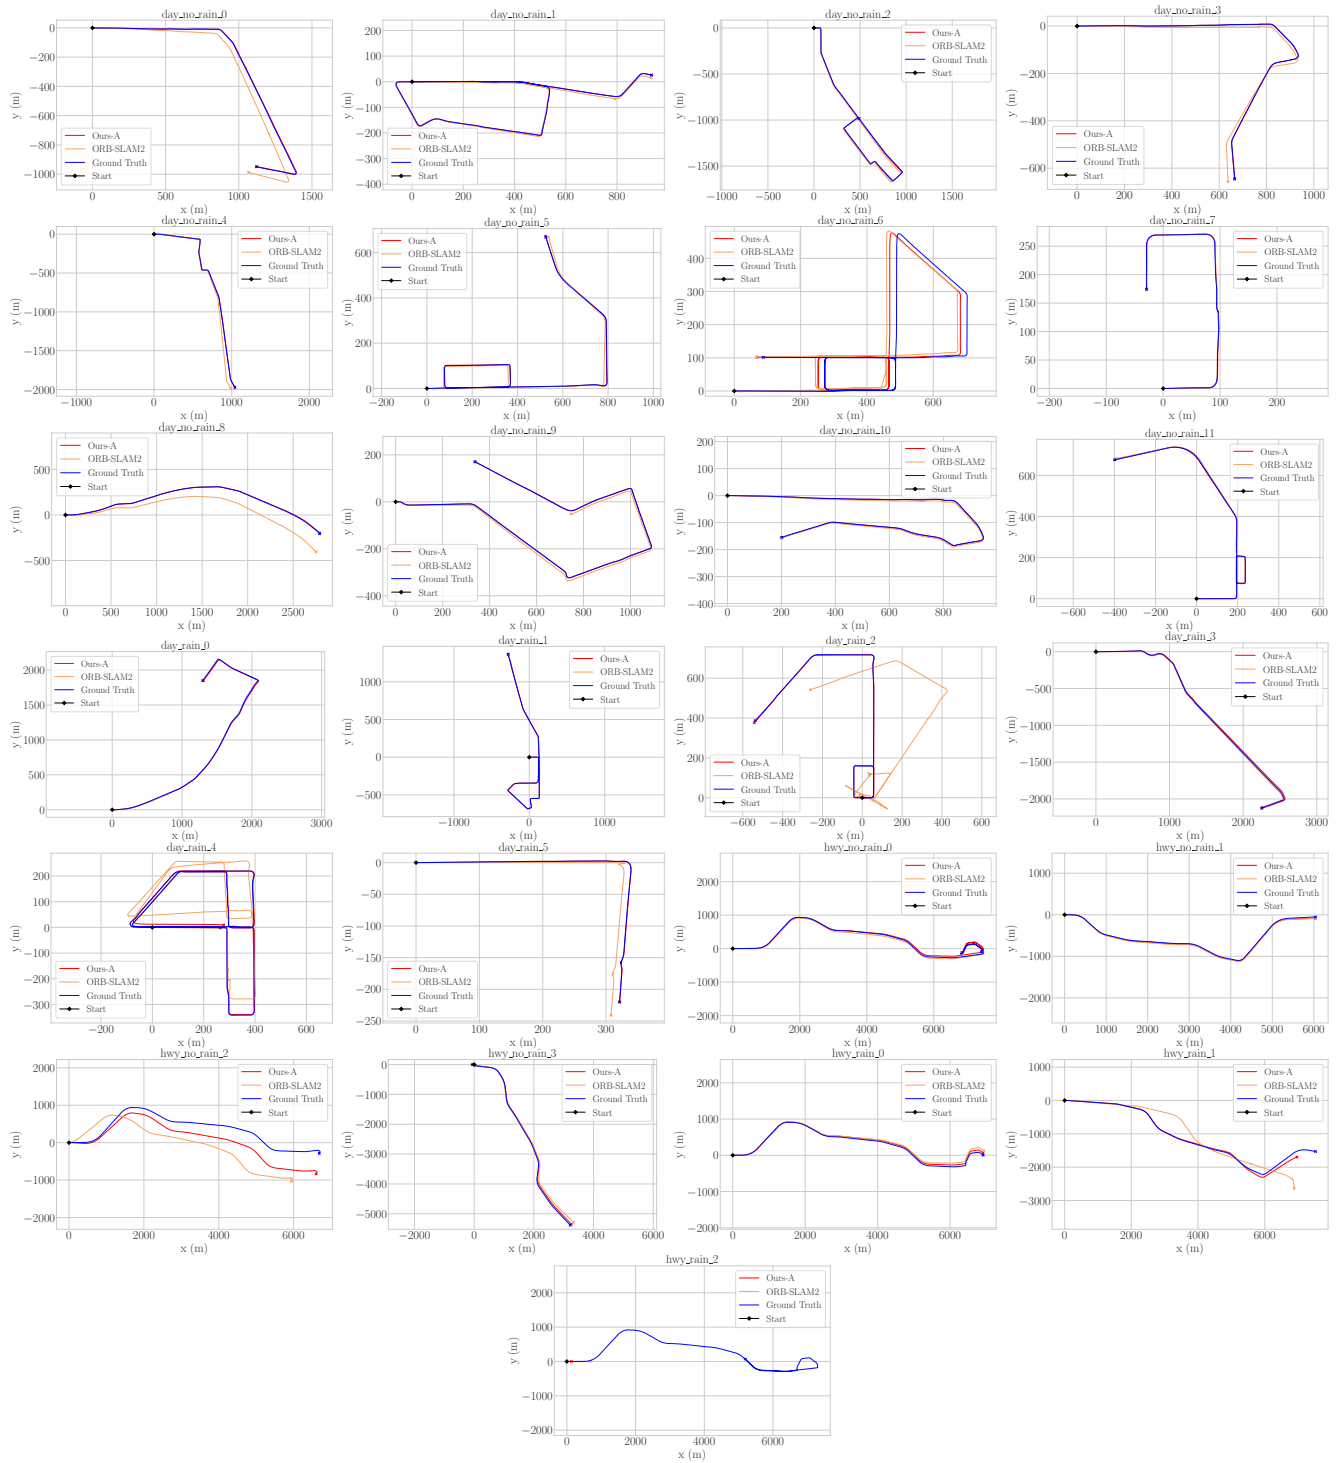
Fig. 13: Estimated trajectories in all 25 validation sequences, comparing ORB-SLAM2 (stereo) with our full system with all 7 cameras.
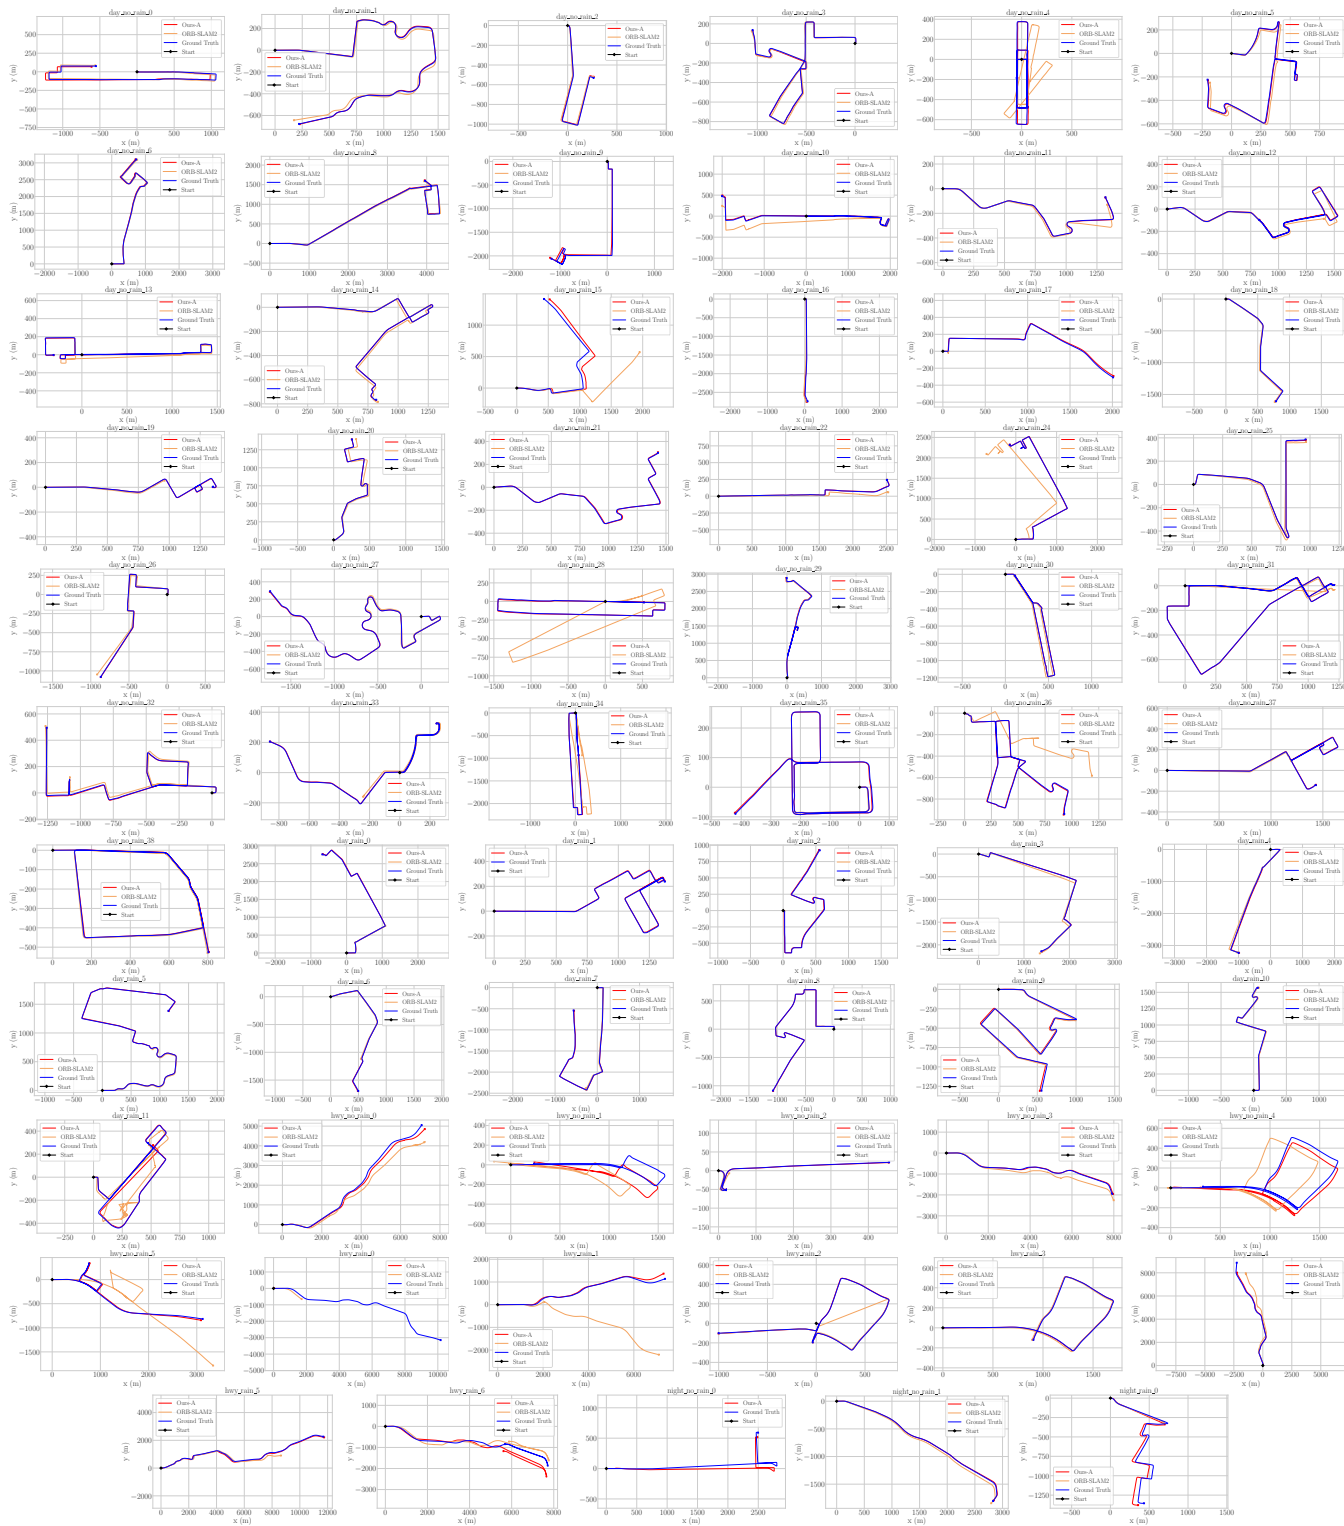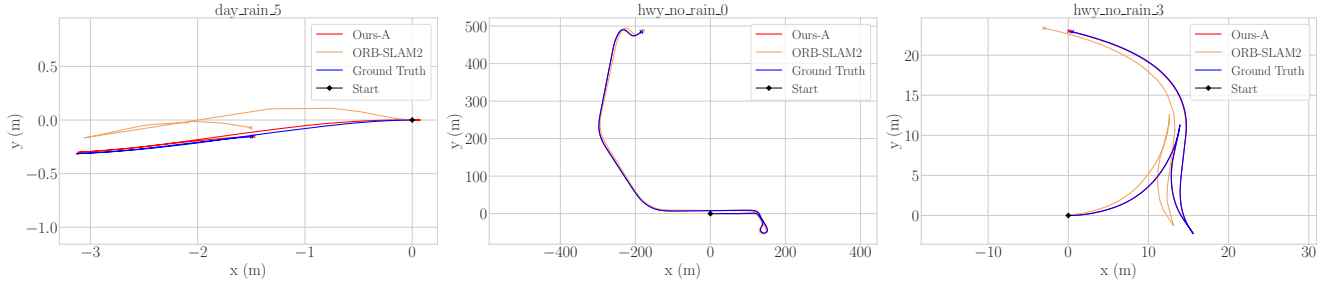
Fig. 14: Estimated trajectories in all 65 train sequences, comparing ORB-SLAM2 (stereo) with our full system (all seven cameras).

Fig. 15: Example maneuvers in the validation set, comparing ORB-SLAM2 and our full system with all 7 cameras. (Left) Reversing into a parallel parking spot. (Middle & Right) Maneuvers in a parking lot.
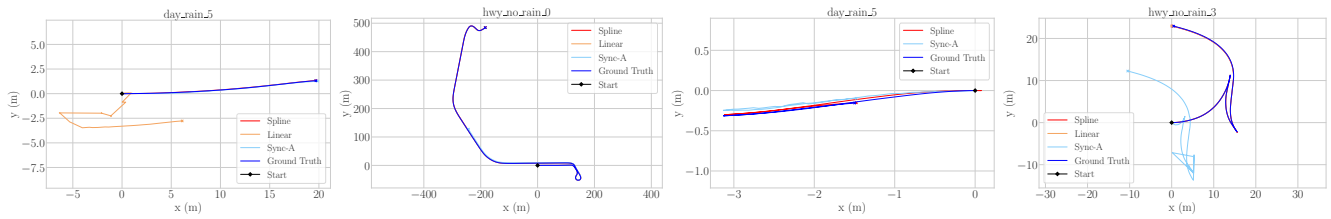
Fig. 16: Motion model ablation trajectories comparing our asynchronous cubic B-spline model, an asynchronous linear motion model, and a discrete-time motion model falsely assuming all cameras are synchronous. (Leftmost) Zoomed-in view on a segment where the linear motion model failed. The vehicle was at an intersection with many dynamic objects. (Right) Maneuvers in the validation set. The linear motion model is missing in the middle sequence because it failed before reaching the maneuver.
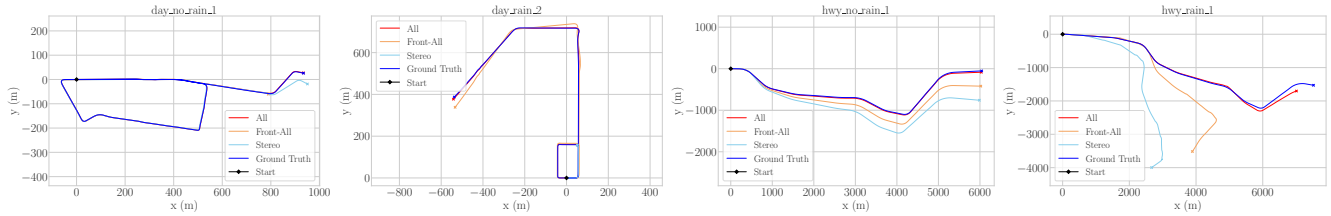
Fig. 17: Camera ablation trajectories estimated with (1) all 7 cameras, (2) all 3 wide front cameras + the stereo pair, and (3) the stereo pair only. The leftmost scenario happened at an intersection where the front view was obstructed by a huge truck that was making a turn. Front-all failed due to repeatedly inconsistent bundle adjustment results, while stereo persisted with a visible rotation error. The second-left scenario is a rainy dusk environment with high volume of traffic. The two scenarios on the right correspond to fast highway driving.
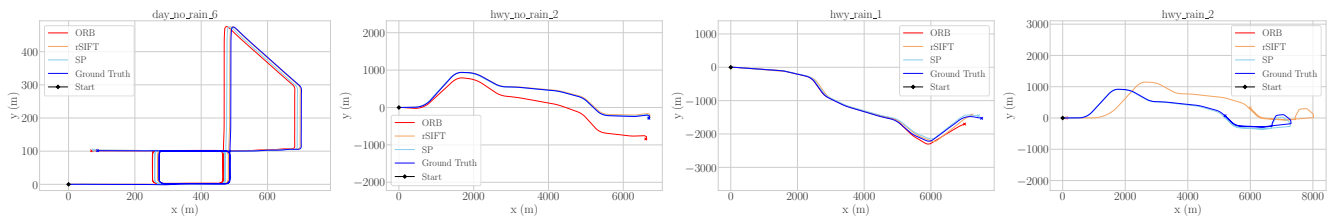

Fig. 18: Keypoint extractor ablation trajectories estimated with ORB, RootSIFT and SuperPoint. RootSIFT and SuperPoint have smaller absolute errors overall and finish a higher percentage of the challenging rainy highway sequences.

Fig. 19: Example where stereo-only loop detection fails due to the presence of the same large bus in two geographically distant frames. This sample is from the training set sequence titled `day_rain_7`.
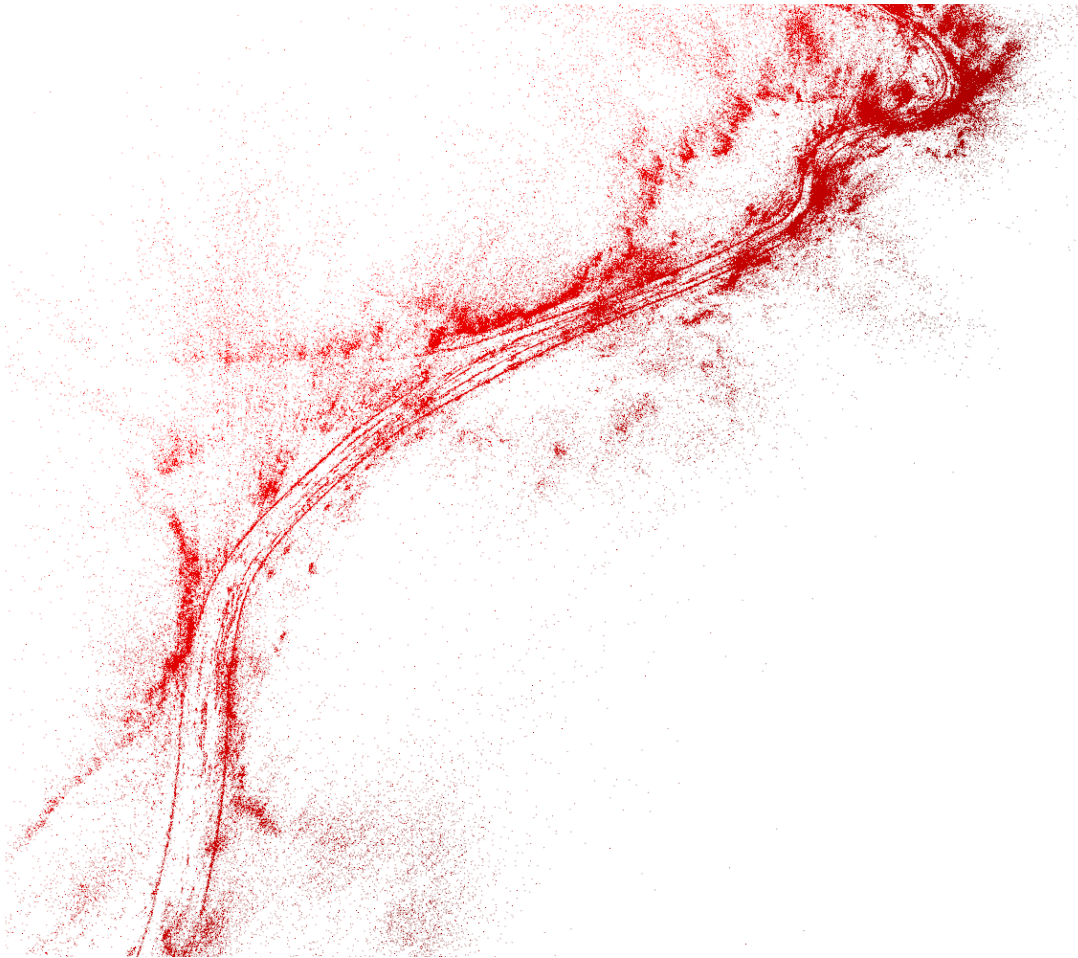

Fig. 20: Qualitative example of the 3D structure produced by the AMV-SLAM system. Note the system's ability to sharply reconstruct the road boundaries, in addition to the surrounding vegetation. This example is from the training set sequence titled `day_rain_5`.
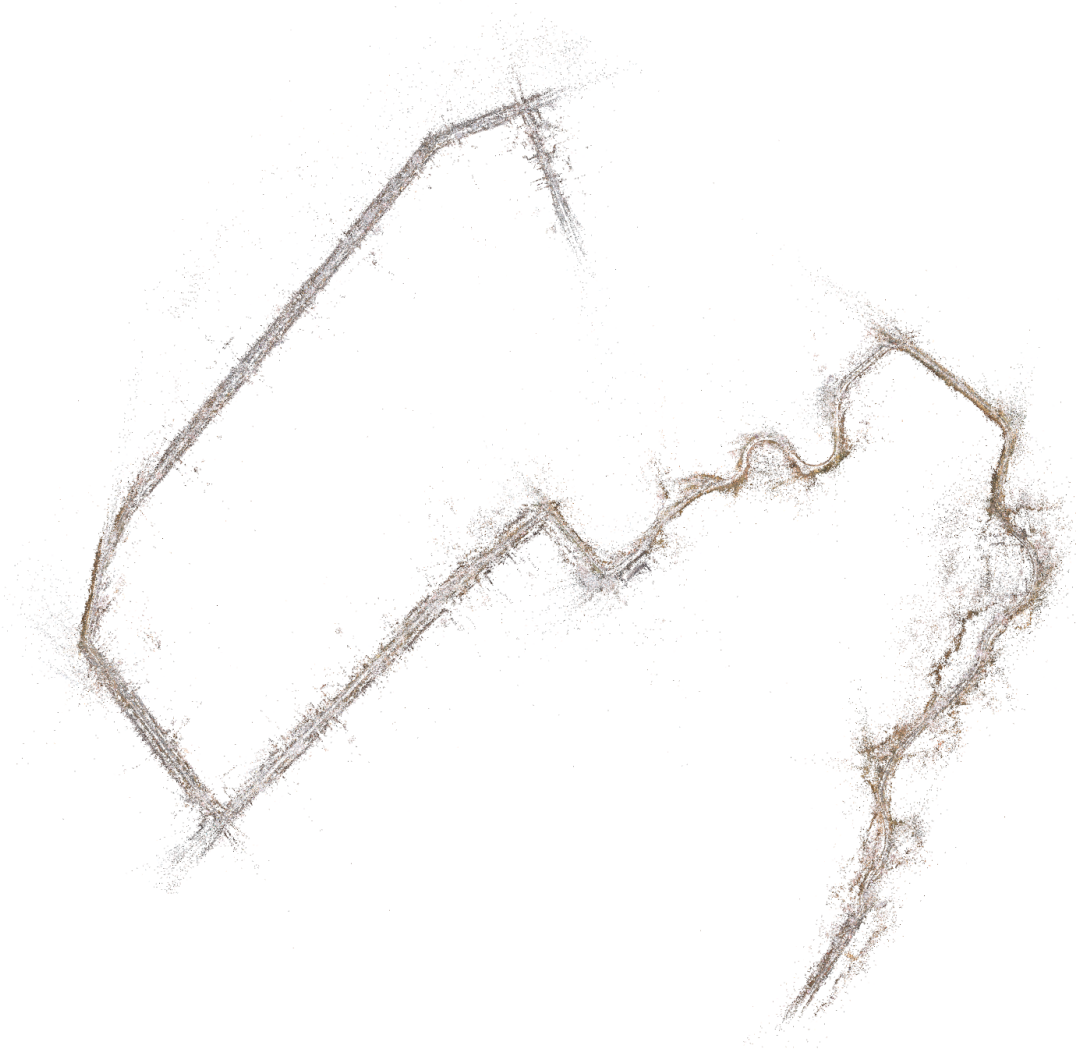
Fig. 21: Reconstructed point cloud from the training sequence `day_rain_5`. Post-processed to include colors from the original camera images. Best viewed in electronic format.
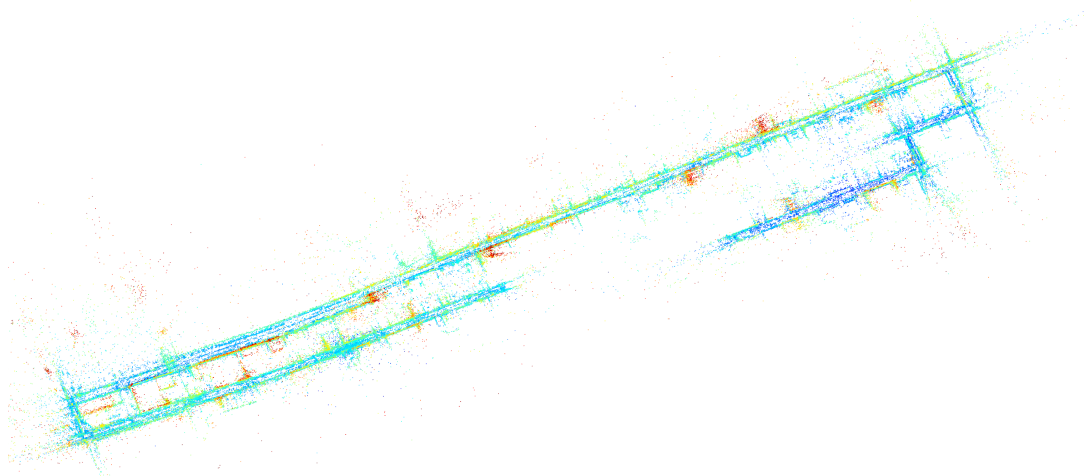

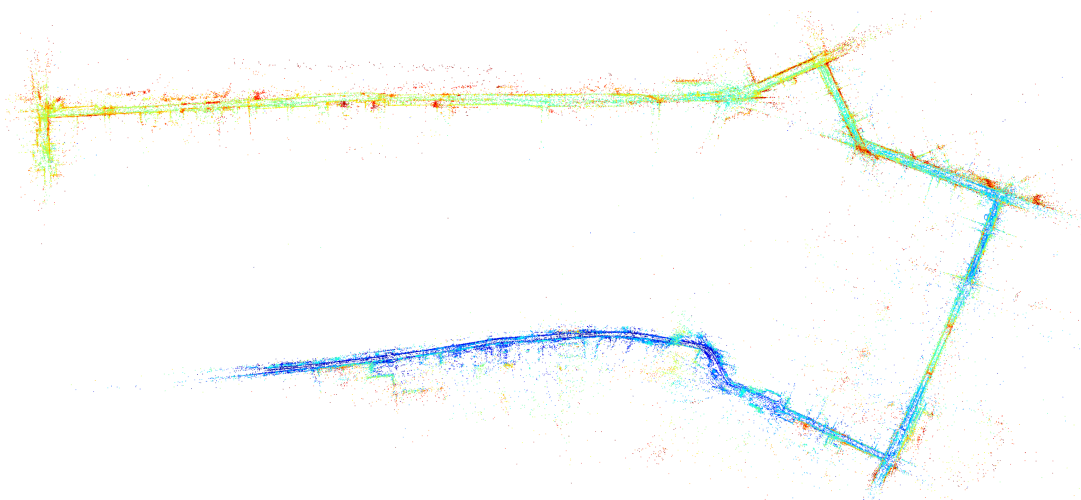Fig. 22: Reconstructed point cloud from the training sequence `day_no_rain_0`.

Fig. 23: Overview of a reconstruction produced from the training set sequence `day_rain_7` by our system. The point cloud is colored by the height (Z) of each point, in the map reference frame.

References

[1] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *ECCV*. Springer, 2014, pp. 834–849. 1, 2

[2] J. Engel, J. Stückler, and D. Cremers, "Large-scale direct SLAM with stereo cameras," in *IROS*. IEEE, 2015. 1, 2

[3] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *PAMI*, vol. 40, no. 3, pp. 611–625, 2017. 1, 2, 6, 13

[4] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, 2015. 1, 2, 3, 4, 5, 6, 8, 9, 11, 13

[5] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: An open-source slam system for monocular, stereo, and RGB-D cameras," *IEEE Trans. Robot.*, vol. 33, no. 5, pp. 1255–1262, 2017. 1, 2, 5, 6, 11, 13, 16

[6] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *IJRR*, vol. 32, no. 11, pp. 1231–1237, 2013. 1, 4, 5, 9

[7] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The EuRoC micro aerial vehicle datasets," *IJRR*, 2016. 1

[8] W. Wang, D. Zhu, X. Wang, Y. Hu, Y. Qiu, C. Wang, Y. Hu, A. Kapoor, and S. Scherer, "TartanAir: A dataset to push the limits of Visual SLAM," Mar. 2020. 1

[9] "Introduction to Intel RealSense visual SLAM and the T265 tracking camera," 2020. 1

[10] J. Xie, M. Kiefel, M.-T. Sun, and A. Geiger, "Semantic instance annotation of street scenes by 3D to 2D label transfer," in *CVPR*, 2016. 1, 5

[11] L. Heng, G. H. Lee, and M. Pollefeys, "Self-calibration and visual SLAM with a multi-camera system on a micro aerial vehicle," in *RSS*, Berkeley, USA, Jul. 2014. 1

[12] M. J. Tribou, A. Harmat, D. W. Wang, I. Sharf, and S. L. Waslander, "Multi-camera parallel tracking and mapping with non-overlapping fields of view," *IJRR*, vol. 34, no. 12, pp. 1480–1500, 2015. 1, 2

[13] A. Harmat, M. Trentini, and I. Sharf, "Multi-camera tracking and mapping for unmanned aerial vehicles in unstructured environments," *Journal of Intelligent & Robotic Systems*, vol. 78, no. 2, pp. 291–317, 2015. 1, 2

[14] S. Urban and S. Hinz, "MultiCol-SLAM - a modular real-time multi-camera SLAM system," *arXiv preprint arXiv:1610.07336*, 2016. 1, 2, 3, 5, 6

[15] P. Liu, M. Geppert, L. Heng, T. Sattler, A. Geiger, and M. Pollefeys, "Towards robust visual odometry with a multi-camera system," in *IROS*, Oct. 2018. 1, 2

[16] "Skydiox2," 2020, (accessed October 9, 2020). [Online]. Available: https://www.skydio.com/pages/skydio-x2 1

[17] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet, "Lyft Level 5 perception dataset 2020," https://level5.lyft.com/dataset/, 2019. 1

[18] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov, "Scalability in perception for autonomous driving: Waymo Open Dataset," in *CVPR*, June 2020. 1, 10

[19] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuScenes: A multimodal dataset for autonomous driving," in *CVPR*, June 2020. 1, 3, 10

[20] Y. Zhou, G. Wan, S. Hou, L. Yu, G. Wang, X. Rui, and S. Song, "DA4AD: End-to-end deep attention aware features aided visual localization for autonomous driving," in *ECCV*, 2020. 1

[21] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part I," *IEEE Robotics and Automation Magazine*, vol. 13, no. 3, pp. 108–117, 2006. 1

[22] J. Leonard, "Directed sonar sensing for mobile robot navigation," Ph.D. dissertation, University of Oxford, 1990. 1

[23] A. Davison and D. Murray, "Mobile robot localisation using active vision," in *ECCV*, 1998. 1

[24] S. Thrun, W. Burgard, and D. Fox, "A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping," in *ICRA*, 2000. 1

[25] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," *AAAI/IAAI*, 2002. 1

[26] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, 2016. 1

[27] R. A. Newcombe, S. J. Lovegrove, and A. J. Davison, "DTAM: Dense tracking and mapping in real-time," in *ICCV*. IEEE, 2011, pp. 2320–2327. 1

[28] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *IJCAI*, 1981, pp. 674–679. 1

[29] M. Irani and P. Anandan, "All about direct methods," in *ICCV Theory and Practice, International Workshop on Vision Algorithms*, 1999. 1

[30] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon, "Bundle adjustment – a modern synthesis," in *International workshop on vision algorithms*. Springer, 1999, pp. 298–372. 1

[31] G. Klein and D. Murray, "Parallel tracking and mapping for small AR workspaces," in *ISMAR*. IEEE Computer Society, 2007, pp. 1–10. 1, 2, 4

[32] P. Wenzel, R. Wang, N. Yang, Q. Cheng, Q. Khan, L. von Stumberg, N. Zeller, and D. Cremers, "4Seasons: A cross-season dataset for multi-weather SLAM in autonomous driving," *arXiv preprint arXiv:2009.06364*, 2020. 2, 5, 10

[33] J. Sola, A. Monin, M. Devy, and T. Vidal-Calleja, "Fusing monocular information in multicamera SLAM," *IEEE Trans. Robot.*, vol. 24, no. 5, pp. 958–968, 2008. 2

[34] G. Hee Lee, F. Faundorfer, and M. Pollefeys, "Motion estimation for self-driving cars with a generalized camera," in *CVPR*, 2013, pp. 2746–2753. 2

[35] X. Meng, W. Gao, and Z. Hu, "Dense RGB-D SLAM with multiple cameras," *Sensors*, vol. 18, no. 7, p. 2118, 2018. 2

[36] C. Zhang, Y. Liu, F. Wang, Y. Xia, and W. Zhang, "VINS-MKF: A tightly-coupled multi-keyframe visual-inertial odometry for accurate and robust state estimation," *Sensors*, vol. 18, no. 11, p. 4036, 2018. 2

[37] W. Ye, R. Zheng, F. Zhang, Z. Ouyang, and Y. Liu, "Robust and efficient vehicles motion estimation with low-cost multi-camera and odometer-gyroscope," in *IROS*. IEEE, 2019, pp. 4490–4496. 2

[38] H. Seok and J. Lim, "ROVINS: Robust omnidirectional visual inertial navigation system," *RA-L*, vol. 5, no. 4, pp. 6225–6232, 2020. 2

[39] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *IROS*, 2013, pp. 1280–1286. 2

[40] P. Furgale, C. H. Tong, T. D. Barfoot, and G. Sibley, "Continuous-time batch trajectory estimation using temporal basis functions," *Int. J. Rob. Res.*, vol. 34, no. 14, p. 1688–1710, Dec. 2015. [Online]. Available: https://doi.org/10.1177/0278364915585860 2

[41] X. R. Li and V. P. Jilkov, "Survey of maneuvering target tracking. part I. dynamic models," *IEEE Transactions on aerospace and electronic systems*, vol. 39, no. 4, pp. 1333–1364, 2003. 2

[42] D. Crouse, "Basic tracking using nonlinear continuous-time dynamic models [tutorial]," *IEEE Aerospace and Electronic Systems Magazine*, vol. 30, no. 2, pp. 4–41, 2015. 2

[43] M. Zefran, "Continuous methods for motion planning," *IRCS Technical Reports Series*, p. 111, 1996. 2

[44] M. Mukadam, X. Yan, and B. Boots, "Gaussian process motion planning," in *ICRA*. IEEE, 2016, pp. 9–15. 2

[45] S. Lovegrove, A. Patron-Perez, and G. Sibley, "Spline Fusion: A continuous-time representation for visual-inertial fusion with application to rolling shutter cameras," in *BMVC*, vol. 2, no. 5, 2013, p. 8. 2, 3, 7

[46] H. Ovrén and P.-E. Forssén, "Trajectory representation and landmark projection for continuous-time structure from motion," *IJRR*, vol. 38, no. 6, pp. 686–701, 2019. 2, 12

[47] J. Hedborg, P. E. Forssen, M. Felsberg, and E. Ringaby, "Rolling shutter bundle adjustment," *CVPR*, pp. 1434–1441, 2012. 2

[48] C. Kerl, J. Stuckler, and D. Cremers, "Dense continuous-time tracking and mapping with rolling shutter RGB-D cameras," in *ICCV*, 2015, pp. 2264–2272. 2

[49] J.-H. Kim, C. Cadena, and I. Reid, "Direct semi-dense SLAM for rolling shutter cameras," in *ICRA*, 2016, pp. 1308–1315. 2

[50] D. Schubert, N. Demmel, V. Usenko, J. Stückler, and D. Cremers, "Direct sparse odometry with rolling shutter," *ECCV*, 2018. 2

[51] T. Schöps, T. Sattler, and M. Pollefeys, "BAD SLAM: Bundle adjusted direct RGB-D SLAM," in *CVPR*, 2019, pp. 134–144. 2

[52] J. Zhang and S. Singh, "LOAM: Lidar odometry and mapping in real-time," in *RSS*, Jul. 2014. 2

[53] H. Alismail, L. D. Baker, and B. Browning, "Continuous trajectory estimation for 3D SLAM from actuated lidar," in *ICRA*, 2014. 2

[54] D. Droeschel and S. Behnke, "Efficient continuous-time SLAM for 3D lidar-based online mapping," in *ICRA*, 2018. 2, 8

[55] J. N. Wong, D. J. Yoon, A. P. Schoellig, and T. Barfoot, "A data-driven motion prior for continuous-time trajectory estimation on SE(3)," *RA-L*, 2020. 2

[56] B. Klingner, D. Martin, and J. Roseborough, "Street view motion-from-structure-from-motion," in *ICCV*, 2013. 2

[57] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urtasun, "End-to-end interpretable neural motion planner," in *CVPR*, 2019, pp. 8660–8669. 2

[58] S. Anderson, F. Dellaert, and T. D. Barfoot, "A hierarchical wavelet decomposition for continuous-time SLAM," in *ICRA*. IEEE, 2014, pp. 373–380. 2

[59] C. Sommer, V. Usenko, D. Schubert, N. Demmel, and D. Cremers, "Efficient derivative computation for cumulative b-splines on lie groups," in *CVPR*, 2020, pp. 11 148–11 156. 2, 3, 12, 14

[60] D. Hug and M. Chli, "On conceptualizing a framework for sensor fusion in continuous-time simultaneous localization and mapping," in *3DV*, 2020. 2

[61] S. Anderson, T. D. Barfoot, C. H. Tong, and S. Särkkä, "Batch nonlinear continuous-time trajectory estimation as exactly sparse Gaussian process regression," *Auton. Robots*, 2015. 2

[62] J. Dong, M. Mukadam, B. Boots, and F. Dellaert, "Sparse Gaussian processes on matrix lie groups: A unified framework for optimizing continuous-time trajectories," *ICRA*, 2018. 2

[63] T. Y. Tang, D. J. Yoon, and T. D. Barfoot, "A white-noise-on-jerk motion prior for continuous-time trajectory estimation on SE(3)," *RA-L*, vol. 4, no. 2, pp. 594–601, 2019. 2

[64] M.-J. Kim, M.-S. Kim, and S. Y. Shin, "A general construction scheme for unit quaternion curves with simple high order derivatives," in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 1995, pp. 369–376. 3, 7

[65] C. De Boor, "On calculating with B-splines," *Journal of Approximation Theory*, vol. 6, no. 1, pp. 50–62, 1972. 3, 7

[66] M. G. Cox, "The numerical evaluation of B-splines," *IMA Journal of Applied Mathematics*, vol. 10, no. 2, pp. 134–149, 1972. 3, 7

[67] S. Agarwal, A. Vora, G. Pandey, W. Williams, H. Kourous, and J. McBride, "Ford Multi-AV seasonal dataset," *arXiv preprint arXiv:2003.07969*, 2020. 3, 5, 10

[68] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A. S. Chung, L. Hauswald, V. H. Pham, M. Mühlegg, S. Dorn *et al.*, "A2D2: Audi autonomous driving dataset," *arXiv preprint arXiv:2004.06320*, 2020. 3, 5, 10

[69] D. Gálvez-López and J. D. Tardós, "Bags of binary words for fast place recognition in image sequences," *IEEE Trans. Robot.*, vol. 28, no. 5, pp. 1188–1197, Oct. 2012. 4, 8

[70] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The Oxford RobotCar dataset," *IJRR*, vol. 36, no. 1, pp. 3–15, 2017. 4, 5, 9

[71] Z. Yan, L. Sun, T. Krajnik, and Y. Ruichek, "EU long-term dataset with multiple sensors for autonomous driving," in *IROS*, 2020. 5

[72] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet large scale visual recognition challenge," *IJCV*, vol. 115, no. 3, pp. 211–252, 2015. 4

[73] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *ECCV*, 2014, pp. 740–755. 4

[74] N. Carlevaris-Bianco, A. K. Ushani, and R. M. Eustice, "University of Michigan North Campus long-term vision and lidar dataset," *IJRR*, vol. 35, no. 9, pp. 1023–1035, 2016. 5, 9

[75] X. Gao, R. Wang, N. Demmel, and D. Cremers, "LDSO: Direct sparse odometry with loop closure," in *IROS*. IEEE, 2018, pp. 2198–2204. 5, 6, 11, 13

[76] E. Rublee, V. Rabaud, K. Konolige, and G. R. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *ICCV*, vol. 11, no. 1, 2011, p. 2. 5, 6, 11, 14, 15

[77] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *IJCV*, vol. 60, no. 2, pp. 91–110, Nov. 2004. 5, 6, 11

[78] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems." in *IROS*, 2012, pp. 573–580. 5

[79] D. DeTone, T. Malisiewicz, and A. Rabinovich, "SuperPoint: Self-supervised interest point detection and description," in *CVPR Workshops*, Jun. 2018. 6, 11, 12, 14, 15

[80] Y. Ono, E. Trulls, P. Fua, and K. M. Yi, "LF-Net: learning local features from images," in *NIPS*, 2018, pp. 6234–6244. 6

[81] M. Dusmanu, I. Rocco, T. Pajdla, M. Pollefeys, J. Sivic, A. Torii, and T. Sattler, "D2-Net: A trainable CNN for joint description and detection of local features," in *CVPR*, 2019, pp. 8092–8101. 6

[82] X. Shen, C. Wang, X. Li, Z. Yu, J. Li, C. Wen, M. Cheng, and Z. He, "RF-Net: An end-to-end image matching network based on receptive field," in *CVPR*, Jun. 2019. 6

[83] J. Revaud, C. De Souza, M. Humenberger, and P. Weinzaepfel, "R2D2: Reliable and repeatable detector and descriptor," in *NIPS*, 2019, pp. 12 405–12 415. 6, 11, 15

[84] R. Arandjelović and A. Zisserman, "Three things everyone should know to improve object retrieval," in *CVPR*, 2012, pp. 2911–2918. 6, 11, 15

[85] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981. 8

[86] S. Agarwal, K. Mierle, and Others, "Ceres solver," http://ceres-solver.org. 11

[87] A. Haarbach, T. Birdal, and S. Ilic, "Survey of higher order rigid body motion interpolation methods for keyframe animation and continuous-time trajectory estimation," in *2018 International Conference on 3D Vision (3DV)*, 2018, pp. 381–389. 12

[88] H. Ovrén and P.-E. Forssén, "Spline error weighting for robust visual-inertial fusion," in *CVPR*, 2018, pp. 321–329. 12