

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

Design and Evaluation of a Generic Visual SLAM Framework for Multi Camera Systems

Pushyami Kaveti¹, Shankara Narayanan Vaidyanathan², Arvind Thamil Chelvan¹, Hanumant Singh¹

Abstract—Multi-camera systems have been shown to improve the accuracy and robustness of SLAM estimates, yet state-of-the-art SLAM systems predominantly support monocular or stereo setups. This paper presents a generic sparse visual SLAM framework capable of running on any number of cameras and in any arrangement. Our SLAM system uses the generalized camera model, which allows us to represent an arbitrary multi-camera system as a single imaging device. Additionally, it takes advantage of the overlapping fields of view (FoV) by extracting cross-matched features across cameras in the rig. This limits the linear rise in the number of features with the number of cameras and keeps the computational load in check while enabling an accurate representation of the scene. We evaluate our method in terms of accuracy, robustness, and run time on indoor and outdoor datasets that include challenging real-world scenarios such as narrow corridors, featureless spaces, and dynamic objects. We show that our system can adapt to different camera configurations and allows real-time execution for typical robotic applications. Finally, we benchmark the impact of the critical design parameters - the number of cameras and the overlap between their FoV that define the camera configuration for SLAM. All our software and datasets are freely available for further research.

Index Terms—SLAM, Field Robots, Data Sets for SLAM

I. INTRODUCTION

SIMULTANEOUS Localization and Mapping (SLAM) is a fundamental task required for the autonomous navigation of mobile robots. Visual SLAM (VSLAM) [1] is one of the more favored approaches as cameras are inexpensive, easily deployable, low-power sensors that provide rich information about the world in which the robot operates. Most VSLAM pipelines [1][2][3][4] use monocular cameras prone to a single point of failure due to degenerate motion, illumination changes, featureless surfaces, and dynamic objects [5]. They are further limited by a small field of view and inconsistencies associated with scale [6][7]. Additional sensors such as Inertial Measurement Units (IMUs) can be used to improve performance in these scenarios, but they can cause large accumulated errors when visual tracking fails. On the other hand, multi-camera systems strike a good balance between data quality and

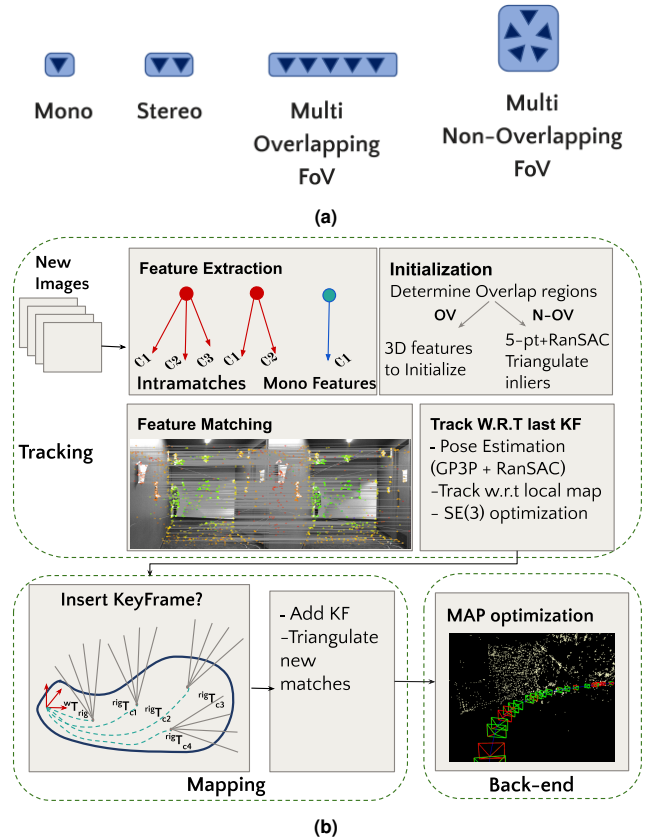


Fig. 1: (a) Illustration of various overlapping(OV) and non-overlapping(N-OV) camera configurations evaluated in this work. (b) Block diagram of the generic visual SLAM framework showing its sub-modules. Feature extraction computes two types of features—multi-view intra-matches and regular mono features. Note the changes, made to initialization, tracking and keyframe representation to adapt to a general multi-camera system.

cost. They can capture more visual information and offer the advantage of redundancy, which helps improve robustness for localization and mapping. As a result, multi-camera sensing has attracted much research interest in recent times leading to novel SLAM solutions [8][9][10], datasets[11], and implementations on robotic platforms[12].

Many earlier works on multi-camera SLAM either treat cameras independently[8][13][14] or assume a specific setup[15][16][17][18], and do not exploit the camera arrangement to the fullest. In this paper, we examine the much more general case of multiple overlapping and non-overlapping cameras. We use a generalized camera model[19] to represent the multi-camera system as a collection of unconstrained rays without making any assumptions about a particular geometry. We build upon the ideas presented in [10] but differ sig-

Manuscript received: June, 23, 2023; Accepted August, 29, 2023.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers' comments.)

¹Pushyami Kaveti, Arvind Thamil Chelvan and Hanumant Singh are with College of Engineering, Northeastern University, Boston, Massachusetts, United States of America (email: {kaveti.p, thamilchelvan.a, ha.singh}@northeastern.edu)

²Shankara Narayanan Vaidyanathan is with Khoury College of Computer Sciences, Northeastern University, Boston, Massachusetts, United States of America (email: vaidyanathan.sh@northeastern.edu)

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

nificantly in terms of modeling the camera system, making our approach extensible to multiple overlapping views beyond stereo pairs. Another challenge with multiple cameras is to effectively and efficiently utilize the increased amount of information provided by the sensors. We compute overlapping regions among cameras in the rig and extract 3D features via cross-matching. This leverages the camera configuration to fuse multi-view data, avoids duplicating features, and keeps the computational cost in check.

Realizing practical mobile robotics applications requires designing sensor systems along with algorithmic development. The choice of cameras and the system configuration of the sensing platform impacts the SLAM outcome[20][21]. While there have been significant efforts toward developing SLAM algorithms for different sensors, there has been a lack of attention to the system design aspects of the sensing platform. To address this gap, we identify the number of cameras and the overlap between their FoVs as two key design parameters of the system configuration that impact the information gathered by the cameras and, thus, SLAM estimates. We evaluate the localization accuracy and robustness of our SLAM system on several real-world datasets collected using our custom-built multi-camera rig. The major contributions of this paper include

- A real-time generic visual SLAM framework that models a multi-camera system as a generalized camera and computes novel 3D intra-match features to fully utilize the overlapping views with open source code¹.
- Evaluation of the SLAM method in terms of tracking accuracy, robustness, and computational constraints.
- A systematic empirical evaluation of different multi-camera system configurations through experiments on real-world datasets in terms of the number of cameras and overlapping vs non-overlapping camera geometries.
- A collection of new real-world datasets² used to evaluate the proposed SLAM system. These datasets allow us to explore more realistic settings with longer runs and far more varying dynamic content than extant SLAM datasets.

II. RELATED WORK

There is a large body of literature on visual SLAM that focuses on monocular, stereo or RGBD cameras and exploits their intrinsic and extrinsic properties. For instance, stereo and RGBD cameras are used to obtain metric SLAM estimates[6][7] in contrast to monocular SLAM[1] with bearing-only image data. Within the monocular realm, wide field-of-view (FoV) and omnidirectional cameras have been used for SLAM[22]. There are many state-of-the-art SLAM pipelines that work off the shelf[2][4].

Multi-camera SLAM has recently gained attention due to its robustness and superior perception capabilities. One of the first multi-camera SLAM frameworks was discussed by Sola et al[8], where data from multiple independent monocular cameras are fused via filtering. Kaess et al[9] propose a probabilistic approach to iteratively solve data association and

SLAM for a circular 8-camera rig. They show that allowing cameras to face different directions yields better constraints and helps with loop closures irrespective of robot orientation. MultiCoSLAM [13] extends ORBSLAM to a rigidly coupled fish-eye multi-camera system by modeling the camera cluster as a generalized camera [19]. Zhang et al[14] present a SLAM system that also calibrates extrinsic parameters using wheel odometry. These works mainly target non-overlapping camera setups that provide maximal coverage of the surroundings and do not exploit the overlap between cameras. Other similar contributions include [16][15][23].

A few works employ stereo pairs to obtain 3D features via stereo matching between overlapping areas and recover metric scale. Heng et al[17] use four fish-eye cameras arranged as two stereo pairs facing forward and backwards to give a 360° view. ROVO[18] uses the same camera configuration but applies a hybrid projection model to minimize distortion and maximize feature matching. In [24], a multi-stereo pipeline is presented with a new 1-point RANSAC for joint outlier rejection. In [25] and its follow-up work[12], multiple stereo pairs improve the feature space for tracking while driving at night from several real-world experiments. In[26], a linear front-facing camera array is used to perform Light Field based static background reconstruction to handle dynamic objects. Li et al [27], perform robust initialization and extrinsic calibration in multi-camera systems with limited view overlaps. Zhang et al[28] use submatrix feature selection to reduce compute and track features across stereo and mono cameras. However, these methods assume a priori knowledge about the camera configuration. Kuo et al[10], present an adaptive SLAM design to accommodate arbitrary camera systems. However, it is not clear if the system can handle overlapping configurations beyond stereo since the method focuses on finding stereo pairs and all the experiments were evaluated on stereo pairs alone.

Studying the impact of camera system design will enable us to build robust multi-camera systems for mobile robots. Nevertheless, the research in this direction is still in a nascent stage. Zhang et al[20] discuss the impact of FoV of a monocular camera on visual odometry in urban environments. Kevin et al[21] benchmark commercially available monocular, stereo and RGB-D sensors by evaluating them on the OpenVSLAM framework. In [16], the authors evaluate MCPTAM on overlapping and non-overlapping camera clusters and show that overlapping cameras recover accurate scale and cause less drift. We add to these efforts by evaluating the impact of the number of cameras and their extrinsic arrangement on SLAM.

III. PROBLEM SETUP

The primary goal of our VSLAM formulation is to develop a unified framework independent of the camera system configuration, which is lightweight and runs in real-time on mobile robots. A block diagram of the generic visual SLAM pipeline is shown in fig. 1.

Definitions: Let the number of cameras in the sensor rig be $N_c \geq 1$. Each component camera c_p is related to the body frame of the multi-camera rig by a rigid body transformation ${}^B T_p = \{{}^B R_p, {}^B t_p\}$ obtained via extrinsic calibration parameters. The intrinsic parameters of the component cameras are

¹ Code is available at <https://github.com/neufieldrobotics/MultiCamSLAM>

²Data available at <https://tinyurl.com/mwfkjrj8k>

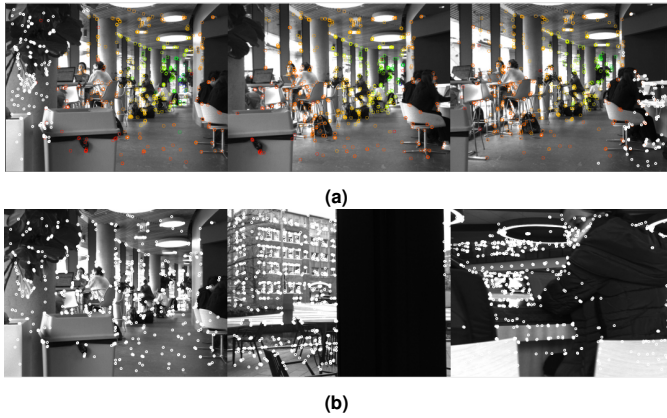


Fig. 2: Two sample multi-camera frames with extracted features on (a) front-facing cameras in overlapping (OV) and, (b) non-overlapping (N-OV) setups for the same scene. The multi-view metric features are colored based on their distance to the camera system. The white points are mono features that do not have any 3D information. Notice that N-OV setup has only mono features, whereas the OV setup has both multi-view and mono features distributed across OV and N-OV regions in the images.

denoted as K_p . Estimated keyframe poses are represented as $x_n \in SE(3)$, $n \in 1, \dots, N$ and correspond to the rig body frame with respect to the world wT_B . The landmarks are represented as $l_m \in R^3$ with $m \in 1, \dots, M$. The visual measurements of the landmarks in the images are denoted as $z_k \in R^2$ with $k \in 1, \dots, K$, the data association that relates each measurement z_k with a keyframe, component camera, and a landmark is given by $(x_{n_k}, c_{p_k}, l_{m_k})$. We assume the body frame of the rig to align with one of the camera coordinate frames for simplicity.

Camera Configuration: An arbitrary camera configuration is the most generic definition, which encompasses all the ways a set of cameras can be rigidly arranged on a robotics platform. For example, we can use monocular or stereo setups, or cameras lying on a circular ring with the possibility of overlapping or non-overlapping FoV. We distinguish different camera configurations based on the number of cameras (N_c) and the FoV overlap among them and benchmark their influence on SLAM. For this study, for the purposes of clarity, we have scoped our work solely on either overlapping (OV) or non-overlapping (N-OV) (shown in fig. 1a) scenarios, even though the methodology is generally applicable to mixing both overlapping and non-overlapping camera rigs. Irrespective of the configuration, the multi-camera rig is considered a single generalized camera that captures a collection of rays passing through multiple pinholes which will be detailed in section IV.

IV. THE FRONT-END

The front end of a SLAM system aims to estimate the pose of the robot and the landmarks observed at each time step. This section discusses the critical aspects of feature extraction and representation, initialization, tracking, and mapping modules that enable us to seamlessly work with arbitrary multi-camera systems.

A. Feature Extraction

We perform sparse feature tracking using two types of features - **Multi-view features** and **Mono features** [Ref. fig. 2].

Depending on their placement, the cameras comprising the camera rig can have overlapping fields of view. We leverage the overlapping imagery to compute strong metric features. Instead of using the features extracted from component cameras independently, we associate overlapping image regions to group features that belong to a specific 3D point in the scene. This differs from most existing camera systems which do not utilize overlap among cameras except to compute stereo. Our multi-view features let us accurately represent the scene with fewer features and avoid creating redundant landmarks.

Determining Overlap: This is the first step required to compute multi-view features efficiently. For each component camera c_i , we find the common image regions in the set of cameras \mathcal{C} that share a common FoV with c_i . Starting with a camera pair (c_i, c_j) , we first divide the image of c_i into a 2D grid. For each grid cell g_k , the epipolar line e_{ij}^k corresponding to its center g_k^x, g_k^y is computed using extrinsic calibration between the camera pair as shown below.

$$F_{ij} = K_j [{}^j t_i]_{\times} {}^j R_i K_i^{-1} \quad (1)$$

$$e_{ij}^k = F_{ij} [g_k^x, g_k^y] \quad (2)$$

Next, we determine the grid cells in the image of c_j that the epipolar line passes through and save them as potential cells to find correspondences. If the line does not pass through the image plane, we know that the camera pair is non-overlapping. This computation is done just once during initialization in the beginning of and then used throughout the run while executing the SLAM framework.

Multi-view Features: These are essentially cross-matches between cameras. First, we extract multi-scale ORB features in all images and assign them to the 2D grid. This process is parallelized for speed. We then iteratively compute feature correspondences across cameras resulting in a total of ${}^{N_c}C_2$ combinations for N_c cameras. For a camera pair (c_i, c_j) , instead of matching every feature in c_i to every feature in c_j , we match features cell-wise based on the overlap to reduce computation. For a set of features, $\mathcal{F}1$, belonging to a cell in the image of c_i we obtain the set of features $\mathcal{F}2$ lying in the cells corresponding to the potential matches obtained from the overlap determination step. We then perform brute-force matching between the sets $\mathcal{F}1$ and $\mathcal{F}2$. Each match is then passed through the epipolar constraint which checks if the corresponding feature in the second view is within a certain distance from the epipolar line. A set of matches \mathcal{M} is created from the first pair of cameras. For the subsequent image pairs, if correspondence is found between two unmatched features, a new match is added to the set of matches \mathcal{M} . If a match is found for an already matched feature, the new feature is added to the existing match.

Each multi-view match represents a 3D feature in the scene. It consists of a bundle of rays given by a set of 2D keypoints, which are projections of the 3D feature in component cameras, the 3D coordinate obtained from triangulation, and a representative descriptor computed to facilitate feature matching for tracking. Note that the 3D point does not need to be observed in all the cameras. The observability depends on the overlap and scene structure.




Fig. 3: Factor graph of the multi-camera back-end with poses X_i , landmarks l_j and the relative camera poses C_p as the variables to be optimized. The black square factor nodes represent constraints on the variables.

Mono Features In the case of either a monocular camera or non-overlapping camera configurations, there are no multi-view matches. Even within overlapping camera configurations, there might be some non-overlapping regions depending on the structure of the 3D scene. We have mono features with a single 2D keypoint and its descriptor to represent the non-overlapping areas.

B. Initialization

This step creates the initial set of landmarks used to track subsequent frames. we perform initialization based on the underlying camera configuration. After extracting features, if the number of metric multi-view features is greater than a certain threshold, we utilize them as the initial map. Otherwise, we must choose two initial frames and compute the relative pose between them. We use a generalized camera model which allows non-central projection suitable to represent multi-camera systems[19]. A key point expressed as a pixel, u , in the image of camera c_p , is represented as a Plucker line $L = [q, q']$ which is a 6-vector made of line direction q and moment q' .

$$\hat{u} = K_p^{-1}u, \quad q = {}^B R_p \hat{u}, \quad q' = {}^B t_p \times q \quad (3)$$

we determine correspondences between two frames and solve the generalized epipolar constraint to obtain relative pose

$$\begin{pmatrix} q_2 \\ q_2' \end{pmatrix}^T \begin{pmatrix} E & R \\ R & 0 \end{pmatrix} \begin{pmatrix} q_1 \\ q_1' \end{pmatrix} = 0 \quad (4)$$

where $[q_1, q_1']$ and $[q_2, q_2']$ are the Plucker rays of matched features, $E = [t]_{\times} R$ is the essential matrix, and R and t are the rotation and translation between the two generalized camera frames. In the case of a monocular camera, the Plucker ray corresponding to the keypoints is $[\hat{u}, 0]$ from eq(3). Thus the generalized epipolar constraint is naturally reduced to the classical epipolar constraint. We apply the 5-point algorithm for a monocular camera and the 17-point algorithm[29] for a multi-camera rig with RANSAC for relative pose estimation. We make sure that there is sufficient parallax between the two frames and triangulate the correspondences which act as our initial landmarks, and follow this with non-linear optimization to obtain our final landmark estimates.

C. Tracking and Mapping

After initialization, every incoming frame is tracked with respect to the last keyframe. We compute inter-frame correspondences between the last keyframe and the current frame

via bag-of-words matching. Since multi-view features contain more than one descriptor from different component cameras we use the median of descriptors for matching. If enough 3D-2D matches are found between the map points ${}^w P_i$ in the last keyframe and observations z_k in the current frame we obtain the Plucker coordinates of z_k , $[q_k, q_k']$ using eq (3) and estimate the absolute pose of the current frame $[{}^B R_w \quad {}^B t_w]$ using generalized PnP[30] by solving a set of constraints of the form

$$\lambda q_i + q_i \times q_i' = {}^B R_w {}^w P_i + {}^B t_w \quad (5)$$

If the estimated pose indicates significant motion since the last keyframe we further localize the current frame with respect to the local map in a manner similar to ORBSLAM [2]. We find the set of neighboring keyframes \mathcal{K} shared by the initially tracked landmarks. We then compute new matches between the landmarks tracked in \mathcal{K} and the current frame. This enables us to gain local map support and helps in finding stable landmarks in the presence of occluding dynamic objects. Finally, the current frame is inserted as a keyframe if the ratio of the tracked landmarks since the last keyframe is less than a certain threshold. When a new keyframe decision is made, the observations are added to the existing landmarks and the new inter-frame matches corresponding to the non-map points are triangulated to create new map points.

V. THE BACK-END

The back end corresponds to the optimization framework that refines the initial estimates of keyframe poses X and landmarks L by maximizing the posterior probability on the variables given observations $Z = \{z_k\}_{k=1}^K$ [31]. In a general multi-camera system an observation not only depends on the pose of the rig X and the landmark L , but also on the component cameras C it is seen in. The Maximum a Posteriori (MAP) problem is given by

$$\begin{aligned} X^*, L^*, C^* &= \operatorname{argmax}_{X, L, C} P(X, L, C | Z) \\ &\propto \operatorname{argmax}_{X, L, C} P(Z | X, L, C) P(X, L, C) \\ &\propto \operatorname{argmax}_{X, L, C} P(x_0) \prod_{k=1}^K P(z_k | x_{n_k}, l_{m_k}, c_{p_k}) \end{aligned} \quad (6)$$

where $P(Z|X, L, C)$ is the likelihood of the observations which factorizes into a product of individual probabilities due to i.i.d assumption. $P(x_0)$ is the prior probability over the initial robot pose. The factor graph representation in fig. 3 shows the individual probability constraints(factors) among variables. Assuming normally distributed zero-mean noise for the observations z_k and modeling prior also as a gaussian eq(6) takes the least-squares minimization form

$$\begin{aligned} X^*, L^*, C^* &\triangleq \operatorname{argmin}_{X, L, C} \sum_{k=1}^K \|h_k(x_{n_k}, l_{m_k}, c_{p_k}) - z_k\|_{\Sigma_k}^2 \\ &\quad - \log p(x_0) \end{aligned}$$

where the measurement function $h_k(x_{i_k}, l_{j_k}, c_{p_k})$ maps a landmark to a predicted observation \hat{z}_k through a series of transformations. First, the pose of the body $x_{n_k} = {}^w T_B^n$ and

Datasets						
Label	Location	Total Frames	Length (m)	Groundtruth	Attributes	Loop
ISEC_Lab1	indoor Lab	9130	152	optitrack	feature less areas , reflective surfaces	Yes
ISEC_Ground1	indoor lobby	4481	90	target based	reflective surfaces, minimal dynamic content, day light	Yes
ISEC_Ground2	indoor lobby	5408	90	target based	reflective surfaces, minimal dynamic content, night	Yes
ISEC_Ground3	indoor lobby	6001	120	target based	reflective surfaces,high dynamic content, day light	Yes
Falmouth	outdoor offroad	9156	2800	target based + gps	fast motion, outdoor, foliage, minimal dynamic content, large area	Yes
Curry_center	campus	28526	590	target based	Urban,outdoor,dynamiccontent,largearea	Yes
ISEC_floor5	campus	4300	187	target based	Urban,indoor,repeated structures	Yes

TABLE I: This table describes the attributes of the dataset used to evaluate the proposed SLAM framework. We collected six sequences with trajectory lengths ranging between 90 m - 2800 m in a wide variety of indoor and outdoor environments.

the relative pose of the component camera $c_{p_k} = {}^B T_{p_k}$ are used to get the pose of the camera in world frame via SE(3) composition ${}^w T_c = {}^w T_B \cdot {}^B T_c$. The 3D landmark $l_{m_k} = P^w$ is transformed from world frame into the camera coordinate frame $P^c = f(T_c^w, P^w) = (R_c^w)^T (P^w - t_c^w)$. Finally, P^c is projected to 2D image coordinates using the intrinsic matrix $\hat{z}_k = K_c P^c$. This formulation conveniently models multi-view features. It gives the back end the flexibility to work with different camera configurations and optimize the extrinsic calibration parameters of the component cameras C along with estimating the trajectory and landmarks.

VI. EXPERIMENTAL SETUP

In this section, we describe the hardware setup, calibration, and datasets collected for the experimental evaluation.

A. Hardware Setup and Camera Calibration

We use a rigid multi-camera rig consisting of 7 cameras with 5 cameras facing forward and 2 sideways and an inertial measurement unit (IMU) as shown in figure fig. 4. The cameras are arranged to accommodate configurations with both overlapping(OV) and non-overlapping(N-OV) FoV as shown in fig. 4. The front-facing cameras(red dashed box) are used to run experiments for mono, stereo, and OV multi-camera setups. The center front-facing and side-facing cameras(blue box) are used as N-OV multi-camera setup. We use FLIR BlackFly 1.3 MP color cameras with a resolution of 720 x 540, 57° FoV and a Vectornav IMU running at 200 HZ. All the cameras are hardware synchronized at 20 fps. We use Kalibr [32] to obtain the intrinsic and extrinsic parameters of the cameras with overlapping FoV. For non-overlapping cameras, target-based calibration does not work as we need the cameras to observe the target to solve for the relative transformation. Instead, we perform imu-camera calibration for each of the non-overlapping cameras and chain the inter-camera transformations together. This gives a good initial estimate of extrinsic parameters, which are refined during optimization as described in section V.

B. Datasets

We collected a set of six indoor and outdoor sequences. The multi-camera rig along with a DELL XPS laptop with 32GB RAM was mounted on a Clearpath Ridgeback robotic




Fig. 4: The custom-built multi-camera rig used to collect data for evaluating the SLAM pipeline. The figure shows overlapping and non-overlapping configurations and the IMU that was mounted on the rig. The IMU is used to compute the baseline between two consecutive cameras is 165mm.

platform and driven across Northeastern University’s campus for data collection. These sequences include several challenging scenes consisting of narrow corridors, featureless spaces, jerky motions, sudden turns, and dynamic objects which are commonly encountered in urban environments. One of the sequences was collected with the NUANCE Autonomous car in an offroad environment. Ground truth is obtained from GPS for the offroad sequence and using Optitrack (accurate up to a millimeter) for the indoor sequence in the lab area. In cases where Optitrack could not be used, visual tags are used as ground truth to compute drift. The details of the dataset including location, length of the trajectory, and ground truth are consolidated in table I.

VII. RESULTS

Here, we discuss the experimental results for several challenging indoor and outdoor trajectories. For quantitative analysis, we use Absolute Translation Error (ATE) obtained by aligning the estimated trajectory with ground truth and computing the mean error between corresponding poses, shown in tables II-IV. When ground truth trajectory is unavailable, we use a visual target at the start and end of the trajectory to compute accumulated drift.

A. Comparing with the State of the Art (SOTA) Algorithms

We compare the performance of our method with existing open source SLAM algorithms.

Stereo SLAM evaluation: We evaluate ORBSLAM3[2] and SVO Pro[4], two popular visual SLAM systems in stereo mode. From table III, we can observe that with a stereo setup, our method exhibits more accuracy in three out of the

	Trajectory length (m)		OV								N-OV					
			2cams		3cams		4cams		5cams		1cam		2cams		3cams	
			ATE (m)	scale	ATE (m)	scale	ATE (m)	scale	ATE (m)	scale	ATE (m)	scale	ATE (m)	scale	ATE (m)	scale
MCSLAM	Plus	11.754	0.0226	1.007	0.022	1.01	0.020	1.01	0.0194	1.01	0.027	3.52	0.023	0.15	0.022	0.204
	Square	9.898	0.011	1.008	0.01	1.007	0.009	1.006	0.009	1.008	0.011	2.01	0.018	0.29	0.0107	0.282
	Circle	11.166	0.039	0.806	0.042	0.807	0.040	0.95	0.043	0.98	0.412	4.42	0.23	0.425	0.086	0.117
MultiColSLAM	Plus	11.754	0.03	1.39	0.16	1.14	0.106	1.248	0.041	1.242	0.021	2.557	0.161	1.431	0.23	0.684
	Square	9.898	0.022	1.522	0.068	0.93	0.142	1.408	0.102	1.646	0.571	1.41	0.54	0.959	0.695	0.578
	Circle	11.166	-	-	-	-	-	-	-	-	-	-	-	-	-	-

TABLE II: Absolute trajectory errors and scale correction across several overlapping (OV) and non-overlapping(N-OV) configurations with respect to MultiColSLAM, the baseline multi-camera SLAM method. For OV setup, we choose cameras from the front-facing array, starting with a stereo setup and adding one camera for each experiment. For the N-OV setup, we choose the 1, 2, and three cameras, respectively, from the set shown in the enclosed blue box in fig. 4 . We outperform the baseline in all sequences. The empty entries correspond to tracking loss. We present the results on three additional simple trajectories with optitrack ground truth, as MultiColSLAM did not work on any datasets shown in fig. 1, which are significantly large and challenging.

Sequence/Algorithm	MCSLAM		ORBSLAM3		SVO Pro	
	ATE (m)	ATE (%)	ATE (m)	ATE (%)	ATE (m)	ATE (%)
ISEC_Lab1	2.65	1.74	2.9	1.91	2.435	1.6
ISEC_Ground1	0.99	1.07	0.85	0.8	0.901	1.01
ISEC_Ground2	2.41	2.51	2.98	3.1	2.698	2.99
Curry_Center	1.18	1.5	4.84	6.3	41.236	11.87
Falmouth	210.02	8.05	49.708	1.91	61.38	2.19
ISEC_floor 5	0.501	0.27	0.516	0.28	0.626	0.33

TABLE III: Absolute translation error(ATE) of the proposed method and two SOTA stereo SLAM systems SVO and ORBSLAM3 with respect to the ground truth for different sequences. The table shows that the stereo configurations for all systems are roughly comparable.

six datasets compared to ORBSLAM3 and SVO. While the error difference is not much for ISEC_Lab1, ISEC_Ground1, and ISEC_Ground2 trajectories, it is significant for the Curry_center trajectory, which is both longer and has much more dynamic content with people moving around. We observed that SVO is more brittle with respect to dynamic content. In the Falmouth sequence, we perform poorly compared to ORBSLAM3 and SVO in the stereo configuration. This sequence, which involves off-road outdoor images, featured dried foliage resulting in poor feature extraction and hence a loss in tracking. This impacted the performance of our method, while the other methods were able to relocalize more accurately. We disabled loop closing in ORBSLAM3 in the comparisons to compute the accumulated drift accurately. With loop closing, ORBSLAM3's estimates may improve. However, for really long trajectories like the Falmouth sequence, we observed that even with loop closing on, ORBSLAM3 could not recover from the accumulated drift.

Figures 6(a) and 6(b) show the estimated trajectories of ISEC_Ground1 and ISEC_Lab1 sequences. In stereo mode, we perform similarly to ORBSLAM3 and SVO and outperform as we increase cameras. At several places along the trajectory, ORBSLAM3 and SVO show artifacts and tracking errors in the estimated poses due to dynamic features (also shown in fig. 5). We can deal with dynamic objects better as we get more support from features distributed across the FoV and are not limited by just the matched stereo features.

Multi-Camera Baseline: We use MultiColSLAM[13] as the baseline multi-camera SLAM method and evaluate it across various OV and N-OV configurations. During our experiments, MultiColSLAM could not track well and failed on all of our datasets shown in table I, irrespective of the multi-camera setup used. For a fair comparison, we collected an additional set of sequences by tracing simple shapes such as pluses, squares, and circles in a feature-rich static environment with Optitrack ground truth. The evaluation results in table II

show that our method outperforms the baseline in all the sequences. We observed that MultiColSLAM is sensitive to rotational motion and experiences continuous tracking failures for the circle sequence. It works on plus and square sequences but displays significant scale errors even when the cameras overlap. This is because the features are matched, tracked, and mapped independently in each camera without leveraging the camera arrangement to get metric features for scale. Figure 7 shows estimated trajectories for 3-camera OV and N-OV setups and the percentage of features that are observed in more than one component camera under the completely overlapping case.

B. Effect of Camera Configuration

Along with SOTA comparisons, we also evaluate and discuss the effect of the following parameters on the performance of the proposed SLAM pipeline.

1) **Accuracy: Number of Cameras:** Within the overlapping configuration, we evaluate our method by choosing a subset of cameras and increasing the number of cameras for each trial. We start with 2 cameras (stereo) with the smallest baseline and go up to 5 cameras in the front-facing array. TableIV shows that within each sequence, the ATE decreases as the number of overlapping cameras increases. We can see the same trend in the trajectory plots shown in fig. 6(a), fig. 6(b) and fig. 5. The estimated

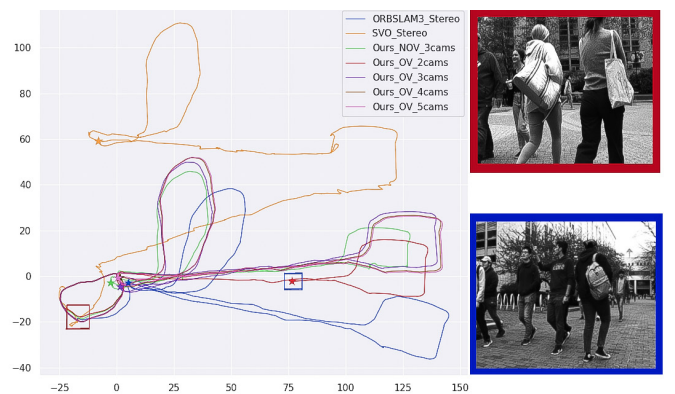


Fig. 5: Estimated trajectories of the Curry_center sequence with outdoor data and dynamic content. Stars indicate final positions of trajectory estimates. Accuracy and robustness improve with increasing number of cameras in OV configurations, as shown by accumulated drift in final position. Red and blue boxes highlight tracking failures caused by occluding dynamic objects. N-OV configuration exhibits scale issues compared to OV configuration but is robust to dynamic content.

	2cams (stereo)		Overlapping configuration						Non-Overlapping configuration			
	ATE (m)	ATE (%)	3cams		4cams		5cams		2cams		3cams	
			ATE (m)	ATE (%)	ATE (m)	ATE (%)	ATE (m)	ATE (%)	ATE (m)	ATE (%)	ATE (m)	ATE (%)
ISEC_Lab1	2.65	1.74	1.23	0.81	0.84	0.56	0.62	0.41	–	–	14.83	9.75
ISEC_Ground1	0.99	1.07	0.350	0.38	0.321	0.34	0.296	0.32	1.13	1.25	1.69	1.88
ISEC_Ground2	2.41	2.51	2.38	2.47	1.61	1.67	1.2	1.25	2.69	2.98	3.20	3.55
ISEC_floor5	0.501	0.27	0.322	0.173	0.317	0.17	0.286	0.153	2.10	1.13	1.774	0.95
Curry_center	1.18	1.5	0.482	0.63	0.359	0.46	0.292	0.38	1.79	2.27	1.423	1.8
Falmouth	210.02	8.05	48.744	1.87	42.179	1.62	28.481	1.09	x	x	x	x

TABLE IV: Analysis of the accuracy of our method for different OV and N-OV configurations. The results indicate a general trend of improved performance as the number of cameras increases. N-OV setups perform poorly compared to their OV counterparts due to scale drift. The symbol ‘-’ indicates tracking loss, while ‘x’ denotes unavailable data.




Fig. 6: Estimated trajectories of the proposed SLAM system, ORBSLAM3, and SVO for ISEC_Ground1 (a) and ISEC_Lab1 (b) sequences. In (b), the ground truth is shown as a dashed line. For ISEC Ground1 sequence (a), the robot’s start and end positions are the same, facilitating performance evaluation. We achieve comparable results to ORBSLAM3 and SVO in stereo setup and demonstrate improved accuracy with increasing overlapping cameras. Notably, ORBSLAM3 and SVO encounter tracking errors due to dynamic objects along the trajectory.

trajectory is closer to ground truth as we increase the cameras as observed in the zoomed part of fig. 6(b).

Overlap vs Non-Overlap: From table IV, errors for the N-OV configuration are greater than the OV configurations in most cases. For instance, the accuracy of the set of 3 overlapping cameras chosen from the front-facing array is higher than the set of the same number of non-overlapping cameras facing different directions. This is because the N-OV setup quickly accumulates scale drift. The error is particularly high in the ISEC_Lab1 sequence, which has narrow featureless corridors and reflective glass walls, rendering the sideward-looking cameras useless for tracking.

2) *Robustness:* To study the robustness of tracking across different camera configurations, we take a closer look at multiple runs of the SLAM pipeline for Curry_center sequence, which is a large dataset (597m) that features heavy dynamic content as shown in fig. 5. This data was collected by driving the robot around campus on a regular working day while there was a lot of human activity. Figure 5 shows the best run for each camera configuration. The red and blue boxes indicate the places where we experience tracking failures, and the corresponding images acquired by one of the cameras are displayed. In our runs, we experience the highest number of tracking failures in the two-camera (stereo) OV configuration. The tracking failures also happen in the 3-camera OV configuration but are less frequent than in the two-camera case. The OV configuration with 4 and 5 cameras runs successfully, closely following each other. The N-OV 3-camera setup does not fail in the presence of dynamic objects because when one view is occluded, it has the support of the other views to track features. However, the trajectory estimates

Average Time Taken for Processing (Milli Seconds)					
	2 cams	3cams	4cams	5cams	N-OV
Feature Extraction	23.1	25.61	29.16	32.64	21.3
Tracking & Mapping	8.02	8.64	9.85	10.78	8.32
Optimization	10.23	20.7	27.02	31.83	10.99
Total Avg time	52.14	64.91	74.89	87.33	58.08

TABLE V: Run time performance of different steps in the SLAM pipeline. The processing time increases with number of cameras. We achieve 11 fps for 5-camera OV configuration which is close to real-time.

are inaccurate and exhibit severe scale errors.

C. Run Time Performance

We conclude the evaluation by measuring the average time taken to process a single multi-camera frame across various camera configurations. Table V reports individual processing times for feature extraction, tracking and mapping, back-end optimization modules, and the total processing time per frame. The results show that the processing time increases with the number of cameras in the OV configuration. This is expected because we have the extra burden of computing multi-view features across cameras in the front end. The computational load in the back end also increases due to increased observations. We can achieve a maximum processing speed of 19.1 fps for a stereo configuration and a minimum of 11.45 fps for five cameras in the OV configuration.

VIII. CONCLUSION

We presented a generic multi-camera SLAM framework that can adapt to any arbitrary camera configuration. The core contribution is the camera configuration independent design and real-time implementation across the complete SLAM pipeline. Leveraging the camera geometry, we extract

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

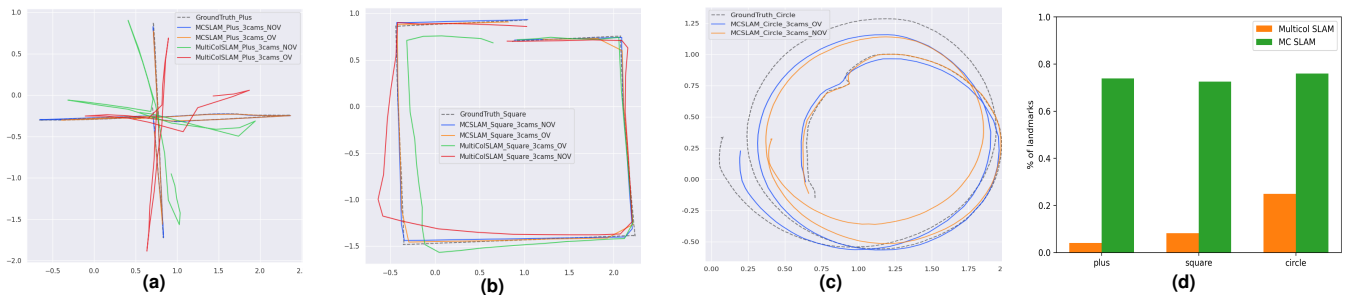


Fig. 7: Performance comparisons between MultiCol SLAM and our system: (a) - (c) show estimated trajectories using a 3-camera setup with OV and N-OV configurations on the simple sequences. (d) Percentage of average no. of landmarks found in more than one component camera by both systems. Our system utilizes metric features more effectively, avoiding scale loss and maximizing feature utilization.

well-distributed multi-view features by effectively using the overlapping FoVs among cameras. We conducted extensive evaluations on real-world datasets collected using a custom-built camera rig featuring a variety of challenging conditions. We also benchmark the performance of the SLAM pipeline in terms of the number of cameras and the overlap information that define the camera configuration. This analysis can be used in designing multi-camera systems for accurate and robust SLAM. This work addresses the gap between the state-of-the-art visual SLAM algorithms and their applicability to real-world deployments of multi-camera systems. We make the code and datasets publicly available to foster research in this direction.

REFERENCES

- [1] A. J. Davison, "Real-time simultaneous localisation and mapping with a single camera," in *Computer Vision, IEEE International Conference on*, vol. 3. IEEE Computer Society, 2003, pp. 1403–1403.
- [2] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial, and Multimodal SLAM," *IEEE Transactions on Robotics*, 2021.
- [3] R. A. Newcombe, S. J. Lovegrove, and A. J. Davison, "DTAM: Dense tracking and mapping in real-time," in *2011 international conference on computer vision*. IEEE, 2011, pp. 2320–2327.
- [4] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "SVO: Semidirect Visual Odometry for Monocular and Multicamera Systems," *IEEE Transactions on Robotics*, vol. 33, pp. 249–265, 2017.
- [5] M. He, C. Zhu, Q. Huang, B. Ren, and J. Liu, "A review of monocular visual odometry," *The Visual Computer*, vol. 36, pp. 1053–1065, 2020.
- [6] J. Engel, J. Stückler, and D. Cremers, "Large-scale direct SLAM with stereo cameras," in *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2015, pp. 1935–1942.
- [7] C. Kerl, J. Sturm, and D. Cremers, "Robust odometry estimation for RGB-D cameras," in *2013 IEEE international conference on robotics and automation*. IEEE, 2013, pp. 3748–3754.
- [8] J. Sola, A. Monin, M. Devy, and T. Vidal-Calleja, "Fusing monocular information in multicamera SLAM," *IEEE transactions on robotics*, vol. 24, no. 5, pp. 958–968, 2008.
- [9] M. Kaess and F. Dellaert, "Probabilistic structure matching for visual SLAM with a multi-camera rig," *Computer Vision and Image Understanding*, vol. 114, no. 2, pp. 286–296, 2010.
- [10] J. Kuo, M. Muglikar, Z. Zhang, and D. Scaramuzza, "Redesigning slam for arbitrary multi-camera systems," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.
- [11] M. Helmberger, K. Morin, N. Kumar, D. Wang, Y. Yue, G. Cioffi, and D. Scaramuzza, "The hilti slam challenge dataset," *arXiv preprint arXiv:2109.11316*, 2021.
- [12] L. Heng, B. Choi, Z. Cui, M. Geppert, S. Hu, B. Kuan, P. Liu, R. Nguyen, Y. C. Yeo, A. Geiger, G. H. Lee, M. Pollefeys, and T. Sattler, "Project autovision: Localization and 3D scene perception for an autonomous vehicle with a multi-camera system," *arXiv*, 2018.
- [13] S. Urban and S. Hinz, "Multicol-slam-a modular real-time multi-camera slam system," *arXiv preprint arXiv:1610.07336*, 2016.
- [14] Z. Zhang and K. Zou, "MMO-SLAM: A Versatile and Accurate Multi Monocular SLAM System," *Journal of Intelligent & Robotic Systems*, vol. 105, no. 3, pp. 1–23, 2022.
- [15] G. Carrera, A. Angeli, and A. J. Davison, "Lightweight SLAM and Navigation with a Multi-Camera Rig," *European Conference on Mobile Robots*, pp. 77–82, 2011.
- [16] 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," *International Journal of Robotics Research*, 2015.
- [17] L. Heng, G. H. Lee, and M. Pollefeys, "Self-calibration and visual slam with a multi-camera system on a micro aerial vehicle," *Autonomous robots*, vol. 39, no. 3, pp. 259–277, 2015.
- [18] H. Seok and J. Lim, "Rovo: Robust omnidirectional visual odometry for wide-baseline wide-fov camera systems," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019.
- [19] R. Pless, "Using many cameras as one," in *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, vol. 2. IEEE, 2003, pp. II–587.
- [20] Z. Zhang, H. Rebecq, C. Forster, and D. Scaramuzza, "Benefit of large field-of-view cameras for visual odometry," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016.
- [21] K. Chappellet, G. Caron, F. Kanehiro, K. Sakurada, and A. Kheddar, "Benchmarking cameras for open vslam indoors," in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021.
- [22] H. Matsuki, L. Von Stumberg, V. Usenko, J. Stückler, and D. Cremers, "Omnidirectional DSO: Direct sparse odometry with fisheye cameras," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, 2018.
- [23] S. Houben, J. Quenzel, N. Krombach, and S. Behnke, "Efficient multi-camera visual-inertial SLAM for micro aerial vehicles," *IEEE International Conference on Intelligent Robots and Systems*, vol. 2016-Novem, no. October, pp. 1616–1622, 2016.
- [24] J. Jaekel, J. G. Mangelson, S. Scherer, and M. Kaess, "A robust multi-stereo visual-inertial odometry pipeline," *IEEE International Conference on Intelligent Robots and Systems*, pp. 4623–4630, 2020.
- [25] P. Liu, M. Geppert, L. Heng, T. Sattler, A. Geiger, and M. Pollefeys, "Towards Robust Visual Odometry with a Multi-Camera System," *IEEE International Conference on Intelligent Robots and Systems*, 2018.
- [26] P. Kaveti, J. S. Nir, and H. Singh, "Towards robust vslam in dynamic environments: A light field approach," in *2021 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2021.
- [27] A. Li, D. Zou, and W. Yu, "Robust initialization of multi-camera slam with limited view overlaps and inaccurate extrinsic calibration," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 3361–3367.
- [28] L. Zhang, D. Wisth, M. Camurri, and M. Fallon, "Balancing the budget: Feature selection and tracking for multi-camera visual-inertial odometry," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, 2021.
- [29] H. Li, R. Hartley, and J.-h. Kim, "A linear approach to motion estimation using generalized camera models," in *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2008, pp. 1–8.
- [30] L. Kneip, H. Li, and Y. Seo, "Upnp: An optimal o(n) solution to the absolute pose problem with universal applicability," in *European conference on computer vision*. Springer, 2014, pp. 127–142.
- [31] F. Dellaert, M. Kaess, et al., "Factor graphs for robot perception," *Foundations and Trends® in Robotics*, vol. 6, no. 1-2, 2017.
- [32] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013.