

Semi-distributed Cross-modal Air-Ground Relative Localization

Weining Lu¹, Deer Bin², Lian Ma², Ming Ma², Zhihao Ma², Xiangyang Chen³,
Longfei Wang³, Yixiao Feng¹, Zhouxian Jiang², Yongliang Shi^{2*}, Bin Liang^{2*}

Abstract—Efficient, accurate, and flexible relative localization is crucial in air-ground collaborative tasks. However, current approaches for robot relative localization are primarily realized in the form of distributed multi-robot SLAM systems with the same sensor configuration, which are tightly coupled with the state estimation of all robots, limiting both flexibility and accuracy. To this end, we fully leverage the high capacity of Unmanned Ground Vehicle (UGV) to integrate multiple sensors, enabling a semi-distributed cross-modal air-ground relative localization framework. In this work, both the UGV and the Unmanned Aerial Vehicle (UAV) independently perform SLAM while extracting deep learning-based keypoints and global descriptors, which decouples the relative localization from the state estimation of all agents. The UGV employs a local Bundle Adjustment (BA) with LiDAR, camera, and an IMU to rapidly obtain accurate relative pose estimates. The BA process adopts sparse keypoint optimization and is divided into two stages: First, optimizing camera poses interpolated from LiDAR-Inertial Odometry (LIO), followed by estimating the relative camera poses between the UGV and UAV. Additionally, we implement an incremental loop closure detection algorithm using deep learning-based descriptors to maintain and retrieve keyframes efficiently. Experimental results demonstrate that our method achieves outstanding performance in both accuracy and efficiency. Unlike traditional multi-robot SLAM approaches that transmit images or point clouds, our method only transmits keypoint pixels and their descriptors, effectively constraining the communication bandwidth under 0.3 Mbps. Codes and data will be publicly available on <https://github.com/Ascbbiac/cross-model-relative-localization.git>.

I. INTRODUCTION

Relative localization is a fundamental problem in heterogeneous multi-robot systems which provides benefits in a wide range of tasks [23], [16], [20]. Current SLAM-based relative localization methods [28], [29], [34], [7] require significant data transmission, making them challenging to work effectively in bandwidth-limited environments. Additionally, their tight coupling with SLAM significantly limits the flexibility of relative localization methods, making them unsuitable for multi-robot systems with heterogeneous sensor configurations, and thereby hindering the execution of downstream tasks.

Leveraging the UAV's extensive field of view and the UGV's superior payload capacity to support additional sensors and computational platforms, our relative localization framework is designed to be decoupled from individual SLAM systems in the multi-robot setting and can be flexibly

integrated as a plugin into air-ground collaborative tasks with different SLAM approaches.

For **efficiency** and **flexibility** of our work, UAV transmits sparse keypoints and corresponding descriptors [22] of its current observation image, along with the global descriptor [1] for place recognition and camera pose (obtained via VIO [10]), to UGV. Similarly, UGV performs the same process on its observation image. However, its camera pose is obtained through linear interpolation of LIO states [5]. Notably, since the map initialization process of the selected LIO, similar to SLAM on UAV, takes the gravity direction into account, the z-axis of the map is always parallel to the gravity direction. Consequently, when estimating the rotational component of the relative pose between aerial and ground systems, only the yaw angle optimization needs to be considered. This reduces the dimensionality of state estimation and improves the efficiency of the estimation process.

For **accuracy**, a hash table [3] of the dense local map is maintained to retrieve the depth of image keypoints from UGV. Once establishing correspondences between keypoints descriptor [22] of UGV and UAV [6], the relative pose is estimated using EPnP [14]. It is evident that this approach alone is prone to significant errors, thus a two-stage optimization is introduced to correct the errors. In the first stage, bundle adjustment of the UGV's local sliding window is employed to correct camera pose errors caused by multi-sensor extrinsic and time offsets, as well as LiDAR depth information corresponding to keypoint pixels. In the second stage, air-ground visual-geometric consistency and relative pose constraint regularization are utilized to optimize the relative pose and camera poses of the UAV.

Furthermore, in the work of the air-ground system, the absence of a common field of view is inevitable. To facilitate long-term relative pose estimation, a database maintenance system of deep learning-based keyframe descriptors is developed, where an incremental Hierarchical Navigable Small Worlds (HNSW) [19], [18] is employed to manage the keyframes on the UGV. These keyframes, represented as global descriptors [1], are derived from both the UAV and UGV keyframes, allowing for the retrieval of co-visibility information from historical frames for the subsequent stage of relative pose estimation.

In summary, our contributions are as follows:

- We propose a semi-distributed cross-modal relative localization architecture for air-ground systems, which decouples the state estimation of individual agents from the overall system.
- A two-stage optimization formulation is proposed to

¹Beijing National Research Center for Information Science and Technology, ²Qiyuan Lab, ³JiangHuai Advanced Technology Center.

* Corresponding author.

Sponsored by the BNRist project (No. BNR2024TD03003) and Qiyuan Innovation Fund (NO.2022-JCJQ-LA-001-107).

achieve efficient and accurate relative state estimation.

- A database maintenance and retrieval system is proposed to build an incrementally updated database utilizing learning-based global descriptors for place recognition.
- Datasets are published proving the superiority in efficiency, accuracy, and flexibility of our method over existing approaches.

II. RELATED WORK

A. Place Recognition

The standard pipeline for localization approaches begins with retrieving prior information, termed place recognition. The robustness, interference resistance, and accuracy of the CNN-based image retrieval algorithm have been substantiated [2]. To enable deployment on mobile devices, a more lightweight visual place recognition approach is proposed [1]. Once the reference images are retrieved, local feature extraction and matching are performed, then the Perspective-n-Point (PnP) algorithm [13] combined with RANSAC is leveraged to estimate the 6-DoF pose. Nonetheless, single-modality approaches exhibit limitations when confronted with challenging scenarios, such as cross-seasonal variations and nighttime conditions. i3dLoc [32] presents an end-to-end visual place retrieval method with respect to offline 3D maps in varying environmental conditions. It retains condition-invariant features by a Generative Adversarial Network (GAN). Tuzcuoglu et al. [27] introduce a cross-modal, cross-view feature matching method between thermal infrared (TIR) and visible images facilitating place retrieval capabilities in nighttime scenarios. Additionally, leveraging higher-level semantic information beyond geometric and textural features enhances both the robustness and discriminative capability of place recognition. Miller utilizes semantics to globally localize a robot using egocentric 3D semantically labelled LiDAR and IMU as well as top-down RGB images obtained from satellites or aerial robots [21]. Gao [9] proposes a two-stage place recognition approach integrating a graph representation of spatial relationship of semantics and class-wise distance fields which encode geometrical cues. LIP-Loc [26] transfers Contrastive Language-Image Pre-Training (CLIP) to the domains of 2D image and 3D LiDAR points to achieve the multi-modal place recognition. However, unlike [8], these approaches are devoid of a comprehensive and systematic framework for database management and retrieval.

B. Cross-modal Relative Localization

Accurate ego-state and relative state estimation are crucial prerequisites for multi-robot missions. Distributed multi-robot SLAM systems with homogeneous sensor setups have been proposed [7], [29], [34], [30], [28]. Dubé et al. [7] construct a centralized Pose Graph Optimization (PGO) utilizing sequential and place recognition constraints for a multiple cooperative robot system. CREPES [30] also leverages a PGO module to integrate the measurements from IR LEDs, cameras, UWBs, and IMUs, achieving highly

accurate and robust relative localization, while it imposes significant requirements for hardware customization. [29], [34], [28] develop decentralized collaborative SLAM systems with different sensor configurations (visual-inertial [28], lidar-inertial [34], omnidirectional visual-inertial-UWB [29]). Despite these advances, the inability to support heterogeneous sensor setups limits system flexibility.

Güler [11] utilizes a leader-follower formation where the aerial robot is equipped with UWB and a camera, while ground robots use only UWB. In SPOMP [20], independent odometry is realized through RGB cameras on the aerial agent and LiDAR on the ground agent, complemented by intermittent semantic information exchange for relative localization. However, they typically integrate relative localization among agents as a critical constraint for their state estimation process.

In contrast, our work is deliberately decoupled from the individual odometry techniques employed by each agent, which enhances the flexibility and interoperability of the system, allowing agents utilizing diverse SLAM algorithms to effectively perform downstream collaborative tasks.

III. METHOD

A. System Overview

As shown in Fig.1, we aim to estimate the state $\mathcal{S} = \{\mathbf{T}_{W_G W_A}, \{\mathbf{T}_{W_A C_A}\}, \{\mathbf{T}_{W_G C_G}\}\}$ and $\mathcal{Z} = \{^{ij}\mathbf{C}\}$, where $\mathbf{T}_{W_G W_A}$ refers to the relative pose in the world coordinate system between the air and ground frames, $\{\mathbf{T}_{W_A C_A}\}$ represents the camera poses in the world coordinate system as measured by the UAV, while $\{\mathbf{T}_{W_G C_G}\}$ denotes the same for the UGV, and $\{^{ij}\mathbf{C}\}$ denotes the depth values corresponding to the j -th visual feature point in frame i at the UGV end. UAV sends frame $^A\mathbf{F}$, including sparse keypoint pixels $^A\mathbf{P} = \{^A\mathbf{p}_1, ^A\mathbf{p}_2, \dots, ^A\mathbf{p}_n\}$, corresponding descriptors $^A\mathbf{D} = \{^A\mathbf{d}_1, ^A\mathbf{d}_2, \dots, ^A\mathbf{d}_n\}$, global descriptor $^A\mathbf{V}$ for place recognition and camera pose $^A\mathbf{T}$. And the UGV will maintain a local sliding window of frame $^G\mathbf{F} = \{\{^G\mathbf{P}\}, \{^G\mathbf{D}\}, ^G\mathbf{T}, ^G\mathbf{V}, \{^G\mathbf{C}\}\}$, of which image is processed in the same manner to obtain keypoints and descriptors $\{^G\mathbf{P}, ^G\mathbf{D}, ^G\mathbf{V}\}$, except that the camera pose $^G\mathbf{T}$ is obtained through linear interpolation of the LIO poses. Furthermore, the points $\{\mathbf{C}\}$ corresponding to the pixels are retrieved from the local point cloud map (Section III-B). On the UGV side, while maintaining its own observation sliding window, it incrementally builds the database with $^A\mathbf{V}$ s acquired from the UAV, facilitating the retrieval of co-visibility information from the UAV's historical observations $^A\mathbf{V}$ s (Section III-C). After obtaining the most similar frames $^A\mathbf{F}$ s of UAV, the UGV first uses its own visual spatiotemporal consistency of $^G\mathbf{F}$ s to correct the camera poses $\{\mathbf{T}_{W_G C_G}\}$ and spatial location of map points $\{\mathbf{C}\}$ (Section III-D.1). Subsequently, optimization of camera poses $\{\mathbf{T}_{W_G C_A}\}$ on UAV is performed through data association between air and ground-based cameras. To ensure numerical stability, the relative pose consistency of different air-ground image pairs is also regarded as a constraint to optimize the relative localization $\mathbf{T}_{W_G W_A}$ (Section III-D.2).




Fig. 1. The system consists of three parts: Data Preprocessing, Database Maintenance and Retrieval, and Two-stage Optimization.

B. Pixel-Point Correspondence

Due to the sparsity of the point cloud of a single LiDAR frame, the keypoint pixels on the UGV index the corresponding depth values from the dense map of LiDAR-SLAM.




Fig. 2. Pixel-Point Correspondence.

Compared to the complete voxelization of octree voxel structures in ray casting [33], iVox [3] offers a sparse voxel representation, establishing and maintaining voxel structures only where spatial points exist. Given the positions of spatial points, a hash table can be employed to uniformly construct and maintain the voxel structure, using the coordinates of the spatial points as keys and the unique indices generated by feeding these coordinates into a hash function as values. As the Fig.2 depicted, given a pixel $G\mathbf{p}$, a ray generated in camera coordinates is iterated along a specified depth range with a fixed step of 0.05 m, starting from shallower depths. At each step, the corresponding points are transformed to world coordinate. Concurrently, in a separate thread, point clouds are accumulated in a hash table, with world coordinates quantized based on a voxel size of 0.1 m. For each converted

point, the nearest neighbor is retrieved from the hash table. If no neighbor is found, the iteration proceeds to deeper points. When a neighbor is identified, the reprojection pixel $G\mathbf{p}'$ is calculated, the reprojection error between $G\mathbf{p}$ and $G\mathbf{p}'$ is evaluated. If the error exceeds a threshold, the iteration continues; otherwise, the nearest neighbor is transformed back to camera coordinates, and the perpendicular foot from the point to the ray serves as the precise 3D representation of the image feature pixel in camera coordinates.

C. Database Maintenance and Retrieval

Although DBoW [8] is widely utilized for incremental place recognition and retrieval in visual SLAM, its reliance on hand-crafted descriptors, such as SURF [4] and ORB [24], renders it incompatible with deep learning-based descriptors, making it unsuitable for the requirements of this study. Given the demonstrated efficiency and accuracy of MixVPR [1] in visual place recognition, coupled with the capability of HNSW [18] to support incremental graph construction, we adopt these methodologies to facilitate keyframe data management and retrieval.

We first quantize the MixVPR into an FP16-format engine, which not only reduces the model's memory but also significantly enhances inference speed. Given the $A\mathbf{V}$ by the MixVPR from the UAV, the HNSW index is constructed incrementally by inserting it into a multi-layered graph. Construction begins at the topmost layer, which has the fewest nodes. For each new $A\mathbf{V}$, the algorithm performs a greedy search from an entry point to identify its nearest neighbors in the current layer. This process iterates downward through all layers up to the element's maximum layer.

In the retrieval process, HNSW employs a multi-layered traversal to navigate the $G\mathbf{V}$ to the most similar $A\mathbf{V}$. The search initiates at the top layer, using a best-first strategy to navigate toward the query's approximate neighbors. This

hierarchical approach reduces the search complexity to approximately logarithmic scale relative to dataset size, as long-range links in higher layers prevent exhaustive traversal.

D. Two-stage Optimization

Air-ground collaborative tasks require execution within a unified global coordinate system. Therefore, the objective of relative localization in this work is to align all agents within the global coordinate system, specifically to determine $T_{W_A W_G}$ (eq. (1)).

$$T_{W_A W_G} = T_{W_A C_A} T_{C_A C_G} T_{W_G C_G}^{-1} \quad (1)$$

Here, $T_{W_A C_A}$ represents the pose of the UAV-mounted camera in the UAV's world coordinate system, $T_{W_G C_G}$ denotes the pose of the UGV-mounted camera in the UGV's world coordinate system, and $T_{C_A C_G}$ indicates the relative pose of the UGV-mounted camera with respect to the UAV-mounted camera. The transformation $T_{W_G C_G}$ is derived via linear interpolation of the LiDAR odometry of the UGV. However, due to the slight variations in the extrinsic parameters between the LiDAR and camera during the movement of the UGV, as well as the lack of precise temporal synchronization between the sensors, certain errors exist in the camera poses obtained through linear interpolation of the LiDAR poses. Additionally, the depths associated with image keypoint pixels are obtained from the map of LiDAR SLAM, and leveraging these points to solve the Perspective-n-Point (PnP) inherently contributes to errors in the estimation of $T_{W_A W_G}$. Additionally, the vio errors inherent to the UAV $T_{W_A C_A}$ can introduce system inaccuracies, leading to instability in $T_{W_A W_G}$. To address the aforementioned issue, we implemented a two-stage optimization process. Specifically, given the key points of the UGV and UAV, along with their corresponding depths on the UGV, we first optimize the camera poses and depths on the UGV. Subsequently, we optimize the relative pose $T_{W_A W_G}$ and the camera poses $\{T_{W_A C_A}\}$ on the UAV simultaneously.

1) **First Stage:** Within a fixed-length sliding window, visual keyframes ${}^G\mathbf{F}_{t:t+n}$, assisted by lightweight and robust feature descriptors $\{{}^G\mathbf{D}\}$ [22]. Therefore, local Bundle Adjustment (BA) [15], [25] is introduced for pose optimization of the keyframes and correction of map point positions. As a gold standard, BA is for the refining structure and initial poses, leveraging the dense correspondences by the [6], it minimizes the reprojection errors:

$$\mathbf{r}_{G_C} = \sum_j \sum_{(i,u) \in \mathcal{T}(j)} \left\| {}^G K ({}^j \mathbf{T}_{W_G C_G}^{-1} \mathbf{C}_i) - {}^G \mathbf{p}_u \right\|_\gamma \quad (2)$$

where $\mathcal{T}(j)$ is the set of frames ${}^G\mathbf{F}$ and keypoints corresponding the local map point index i , $\|\cdot\|_\gamma$ is a robust norm, ${}^G K$ is intrinsic parameters of the camera on UGV, then more accurate $\{T_{W_G C_G}\}$ and $\{C_i\}$ will be refined. Specially, to ensure numerical stability during the optimization process, we optimize the 3D map points using inverse depth parametrization. In doing so, we implicitly correct the errors introduced by extrinsic parameters, time unsynchronization, and linear interpolation.

2) **Second Stage:** Given the refined $\{T_{W_G C_G}\}$ and $\{C_j\}$, they are fixed in the subsequent optimization process. This stage establishes a nonlinear optimization incorporating dual constraint mechanisms, with the objective function formulated as:

$$\mathbf{r} = \mathbf{r}_{C_G C_A} + \mathbf{r}_{W_A W_G} \quad (3)$$

$\mathbf{r}_{C_G C_A}$ denotes the reprojection error term (eq. (4)) derived from air-ground cross-platform dense feature correspondences, which simultaneously integrates spatial geometric constraints and multi-view texture consistency metrics. To get the correspondences, with the method from section III-C, the $\{{}^G\mathbf{F}_{t:t+n}\}$ on the UGV retrieve the most similar UAV frames $\{{}^A\mathbf{F}_{t':t'+n}\}$ from the database, and further filters reliable frames by using FAISS [6] for matching, discarding pairs with fewer matches than a predefined threshold.

$$\mathbf{r}_{C_G C_A} = \sum_j \sum_{(i,u) \in \mathcal{T}(j)} \left\| {}^A K ({}^j \mathbf{T}_{W_A C_A}^{-1} \mathbf{T}_{W_A W_G} \mathbf{T}_{W_G C_G} \mathbf{C}_i) - {}^A \mathbf{p}_u \right\|_\gamma \quad (4)$$

It means that the $\{T_{W_A W_G}\}$ and camera poses of the UAV are both adjusted within the sliding window. Additionally, to enhance the robustness of pose graph optimization, we construct the relative pose constraint regularization term $\mathbf{r}_{W_G W_A}$ (eq. (5)), which enforces compatibility conditions on the Lie manifold for relative poses $\{T_{W_A W_G}\}$ derived from all air-ground keyframe pairs. By minimizing the residual $\mathbf{r}$, we simultaneously adjust $\{T_{W_A C_A}\}$ to optimize and obtain an accurate $T_{W_A W_G}$.

$$\mathbf{r}_{W_A W_G} = \sum_j \left\| \mathbf{T}_{W_A W_G} {}^j \hat{\mathbf{T}}_{W_A W_G}^{-1} - \mathbf{I} \right\| \quad (5)$$

IV. EXPERIMENTS AND ANALYSIS

A. Platform

The experimental platform consists of two main components: UGV and UAV (Fig.3). The UGV utilizes a robust AgileX Scout Mini, equipped with an Azure Kinect camera, an Ouster LiDAR OS-128 for comprehensive environmental perception, and an onboard mini PC featuring an Intel i9-13900H processor alongside an NVIDIA RTX 4070 GPU with 8GB of VRAM. The UAV is outfitted with a Pixhawk6 Pro flight control module for enhanced stability and precision. It integrates a RealSense D435i for visual sensing and a Livox Mid-360 LiDAR to facilitate comparative analysis with LiDAR-based relative positioning methods. The core computing unit of the UAV is a Jetson Orin NX, boasting 16GB of VRAM. In addition, the communication for our air-ground collaboration utilizes a video transmission device with a bandwidth of 40 Mbps. This setup enables rigorous testing and validation of our methodologies in diverse scenarios.

B. Dataset

We provide two datasets corresponding to scenarios with varying levels of co-visibility: sufficient co-visibility *Dataset1* and no co-visibility *Dataset2* at the initial state.



Fig. 3. Hardware Configurations of UAV and UGV.

Each dataset includes synchronized camera, LiDAR, and IMU data, along with pre-calibrated parameters. The calibration data comprises camera intrinsics, as well as extrinsic parameters between the camera, LiDAR, and IMU. To ensure high-quality ground truth, we employ the FZMotion Motion Capture System to record the pose trajectories of both the UAV and UGV during their operations. The presence of comprehensive sensor data and accurate calibration parameters facilitates benchmarking across different approaches, such as single-modal-based and cross-modal-based methods. The inclusion of ground truth trajectories enables precise assessment of reconstruction accuracy and relative pose estimation, making the datasets valuable for advancing research in air-ground collaboration, especially in GPS-denied scenarios.

C. Metrics

We assume that the gravity estimation at both the UAV and UGV ends is accurate, with the z-axis of each map aligned parallel to the gravity direction. Consequently, our relative pose estimation consists of only 4 DoF: x, y, z, and yaw. During quantitative evaluation, translational error is measured using the Euclidean distance (m) in three-dimensional space, while rotational error is assessed solely by the deviation of the yaw angle ($^{\circ}$). The duration of state estimation (ms) serves as a metric for evaluating algorithm efficiency. In comparison with other methods, the stability and robustness of the algorithm are evaluated based on the success rate (SR) of pose estimation. In tables presented in the A/B format, A denotes the number of successful trials, while B represents the total number of tests conducted.

D. Ablation Study

The cost function $r_{W_G W_A}$ avoids the instability of state estimation dependent on a single data pair, though it offers limited improvement in accuracy. $r_{C_G C_A}$ leverages the spatial consistency in the common-view region between the UGV and UAV for local BA to refine the relative pose. r_{G_C} takes into account errors caused by the extrinsic parameters, time offset, and pixel-cloud correspondence between sensors on the UGV side, and corrects these errors through local BA. The results in Table I further substantiate our claims, highlighting the rationality and superiority of our method across different datasets. Nevertheless, as the complexity of the optimization process increases, while the accuracy can reach centimeter or sub-centimeter levels, the computational

TABLE I
EVALUATION FOR STATE ESTIMATION OF DIFFERENT SETTINGS

Dataset	r_{G_C}	$r_{C_G C_A}$	$r_{W_G W_A}$	Translational Error (m)	Rotation Error ($^{\circ}$)	Time (ms)
Dataset1	×	×	✓	0.187	6.42	3.86
	×	✓	✓	0.137	6.04	311
	✓	✓	✓	0.065	5.36	462
Dataset2	×	×	✓	0.129	6.24	5.27
	×	✓	✓	0.090	6.20	307
	✓	✓	✓	0.004	5.07	556

time also increases by more than 100 times. In practical experiments, however, performing relative localization once per second is fully sufficient to meet the requirements, achieving a favorable balance between precision and efficiency.

E. Performance with Different Keypoint Counts

Limited communication bandwidth in air-ground collaborative systems can significantly impact the performance of relative localization. In our system, the UAV transmits per-frame data packets to the UGV containing four principal components: n 2D keypoint coordinates and 64D feature descriptors, 7-DoF pose representation (quaternion + translation), 512D global descriptor. Given 4-byte quantization per dimension, the total bandwidth demand per second is formulated as:

$$W = 4 \times (2n + 64n + 7 + 512) \text{ bytes} \Rightarrow \frac{W \times 8}{10^6} \text{ Mbps} \quad (6)$$

As shown in the Table II, the increase in the number of keypoints leads to a more significant improvement in accuracy, but this also results in higher computation time. A total of 1,024 keypoints is sufficient to achieve sub-centimeter-level accuracy, while requiring a bandwidth of only 2.18 Mbps. In scenarios where lower precision is acceptable, our bandwidth requirement can be further reduced to less than 0.3 Mbps. This demonstrates an efficient trade-off between precision and bandwidth consumption, highlighting the effectiveness of our approach in resource-constrained environments.

F. Evaluation for Database Maintenance and Retrieval

As an alternative to DBow3, our method achieves better compatibility with neural network-based global descriptors. We evaluate both methods across three key metrics: average single-frame incremental insertion time (SIT), including feature extraction and database construction; average retrieval time per single frame (SRT); and accuracy (Acc). Using our actual recorded data, we incrementally constructed a database and retrieved the most similar candidate, as illustrated in Table III. The SIT differs by approximately 1 ms between the two methods, with both ensuring real-time performance. Specifically, our method has an SIT of 26.36 ms compared to DBow3's 25.16 ms. In terms of accuracy, our method exhibits a slight advantage, achieving about a 2% improvement over DBow3, with accuracies of 82% and 80%, respectively. Regarding retrieval efficiency, our method demonstrates a substantially greater advantage, completing

TABLE II
PERFORMANCE WITH DIFFERENT KEYPOINT COUNTS

Keypoint Counts	Dataset1			Dataset2			Bandwidth (Mbps)
	Rotation Error(°)	Trans Error(m)	Time(ms)	Rotation Error(°)	Trans error(m)	Time(ms)	
128	6.60	0.341	115	8.56	0.204	73	0.29
256	6.29	0.305	253	10.97	0.112	182	0.56
512	9.68	0.209	346	8.76	0.052	445	1.10
1024	5.36	0.065	462	5.07	0.004	556	2.18

TABLE III
DBOW3 VS OURS

Methods	SIT (ms)	SRT (ms)	Acc (%)
Ours	26.36	0.84	82
DBow3	25.16	21.14	80

TABLE IV
EVALUATION FOR STATE ESTIMATION WITH DIFFERENT METHODS

Dataset	Methods	Translational Error (m)	Rotation Error (°)	SR	Time (ms)
Dataset1	D ² SLAM	0.554	66.37	2/10	27.2
	FPFH-Teaser	0.039	7.162	6/10	5579
	Ours	0.065	5.361	10/10	465
Dataset2	D ² SLAM	×	×	×	×
	FPFH-Teaser	0.029	5.693	7/10	4963
	Ours	0.004	5.066	10/10	556

retrieval within 0.84 ms—approximately 25 times faster than DBow3, which completes retrieval in 21.14 ms. These results highlight the superior efficiency and effectiveness of our proposed approach.

G. Comparasion with Other Methods

We compared our approach against state-of-the-art baselines. We selected D²SLAM [28], a leading vision-based multi-robot SLAM, as the primary benchmark. We also deployed SlideSLAM [17], a state-of-the-art multi-robot LiDAR SLAM method that heavily relies on semantic information, but it failed to achieve relative localization. Consequently, we chose TEASER [31], one of the SOTA point cloud registration algorithm integrated within hdl_localization [12], for evaluation. As shown in Table IV, D²SLAM relies heavily on sufficient common-view information during initialization. Therefore, it cannot perform relative localization in *Dataset2* without common-view data in the initial state. Even with sufficient common-view data, its performance remains unstable, with limited accuracy. While TEASER achieves favorable results, it requires the pre-computation of FPFH descriptors. Moreover, this approach has specific requirements for point cloud density parameters and map scale, and failure to satisfy these conditions may lead to localization failure. Our method achieves the best balance in terms of stability, accuracy, and efficiency.

V. CONCLUSIONS

In conclusion, this paper presents a novel semi-distributed cross-modal air-ground relative localization framework that decouples from the state estimation of multiple robots, enhancing both flexibility and accuracy in collaborative tasks. By leveraging the UGV’s ability to integrate diverse sensors, the proposed method independently performs SLAM on both the UGV and UAV while utilizing deep learning-based keypoint descriptors and global descriptors for localization. The two-stage optimization process ensures rapid and accurate relative pose estimation, while the database maintenance method ensures incremental and real-time keyframe insertion and retrieval. Experimental results demonstrate significant improvements in localization performance, while the keypoint-based communication strategy effectively reduces bandwidth usage, making the approach suitable for bandwidth-constrained environments. This work provides a promising direction for future air-ground collaboration in GPS-denied and bandwidth-limited scenarios, with potential applications in both wilderness and post-disaster settings.

REFERENCES

- [1] Amar Ali-Bey, Brahim Chaib-Draa, and Philippe Giguere. Mixvpr: Feature mixing for visual place recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2998–3007, 2023.
- [2] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016.
- [3] Chung-e Bai, Tao Xiao, Yajie Chen, Haoqian Wang, Fang Zhang, and Xiang Gao. Faster-lid: Lightweight tightly coupled lidar-inertial odometry using parallel sparse incremental voxels. *IEEE Robotics and Automation Letters*, 7(2):4861–4868, 2022.
- [4] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 110(3):346–359, 2008.
- [5] Kenny Chen, Ryan Nemiroff, and Brett T Lopez. Direct lidar-inertial odometry: Lightweight lid with continuous-time motion correction. In *2023 IEEE international conference on robotics and automation (ICRA)*, pages 3983–3989. IEEE, 2023.
- [6] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library. *arXiv preprint arXiv:2401.08281*, 2024.
- [7] Renaud Dubé, Abel Gawel, Hannes Sommer, Juan Nieto, Roland Siegwart, and Cesar Cadena. An online multi-robot slam system for 3d lidars. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1004–1011. IEEE, 2017.
- [8] Dorian Gálvez-López and Juan D Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on robotics*, 28(5):1188–1197, 2012.

- [9] Peng Gao, Jing Liang, Yu Shen, Sanghyun Son, and Ming C Lin. Visual, spatial, geometric-preserved place recognition for cross-view and cross-modal collaborative perception. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11079–11086. IEEE, 2023.
- [10] Patrick Geneva, Kevin Eickenhoff, Woosik Lee, Yulin Yang, and Guoquan Huang. Openvins: A research platform for visual-inertial estimation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4666–4672. IEEE, 2020.
- [11] Samet Güler and İsa E. Yıldırım. A distributed relative localization approach for air-ground robot formations with onboard sensing. *Control Engineering Practice*, 135:105492, 2023.
- [12] Kenji Koide, Jun Miura, and Emanuele Menegatti. A portable three-dimensional lidar-based system for long-term and wide-area people behavior measurement. *International Journal of Advanced Robotic Systems*, 16(2):1729881419841532, 2019.
- [13] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Real-time solution to the absolute pose problem with unknown radial distortion and focal length. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2816–2823, 2013.
- [14] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Ep n p: An accurate o(n) solution to the p n p problem. *International journal of computer vision*, 81:155–166, 2009.
- [15] Philipp Lindenberger, Paul-Edouard Sarlin, Viktor Larsson, and Marc Pollefeys. Pixel-perfect structure-from-motion with featuremetric refinement. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5987–5997, 2021.
- [16] Björn Lindqvist, Samuel Karlsson, Anton Koval, Ilias Tevetzidis, Jakub Haluška, Christoforos Kanellakis, Ali akbar Agha-mohammadi, and George Nikolakopoulos. Multimodality robotic systems: Integrated combined legged-aerial mobility for subterranean search-and-rescue. *Robotics and Autonomous Systems*, 154:104134, 2022.
- [17] Xu Liu, Jiuzhou Lei, Ankit Prabhu, Yuezhan Tao, Igor Spasojevic, Pratik Chaudhari, Nikolay Atanasov, and Vijay Kumar. Slideslam: Sparse, lightweight, decentralized metric-semantic slam for multi-robot navigation. *arXiv preprint arXiv:2406.17249*, 2024.
- [18] Yu A. Malkov and D. A. Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(4):824–836, 2020.
- [19] Yury Malkov, Alexander Ponomarenko, Andrey Logvinov, and Vladimir Krylov. Approximate nearest neighbor algorithm based on navigable small world graphs. *Information Systems*, 45:61–68, 2014.
- [20] Ian D. Miller, Fernando Cladera, Trey Smith, Camillo Jose Taylor, and Vijay Kumar. Air-ground collaboration with spomp: Semantic panoramic online mapping and planning. *IEEE Transactions on Field Robotics*, 1:93–112, 2024.
- [21] Ian D. Miller, Anthony Cowley, Ravi Konkimalla, Shreyas S. Shivakumar, Ty Nguyen, Trey Smith, Camillo Jose Taylor, and Vijay Kumar. Any way you look at it: Semantic crossview localization and mapping with lidar. *IEEE Robotics and Automation Letters*, 6(2):2397–2404, 2021.
- [22] Guilherme Potje, Felipe Cadar, André Araujo, Renato Martins, and Erickson R Nascimento. Xfeat: Accelerated features for lightweight image matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2682–2691, 2024.
- [23] Yara Rizk, Mariette Awad, and Edward W Tunstel. Cooperative heterogeneous multi-robot systems: A survey. *ACM Computing Surveys (CSUR)*, 52(2):1–31, 2019.
- [24] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. Ieee, 2011.
- [25] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [26] Sai Shubodh, Mohammad Omama, Husain Zaidi, Udit Singh Parihar, and Madhava Krishna. Lip-loc: Lidar image pretraining for cross-modal localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 948–957, 2024.
- [27] Önder Tuzcuoğlu, Aybora Köksal, Buğra Sofu, Sinan Kalkan, and A Aydın Alatan. Xoftr: Cross-modal feature matching transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4275–4286, 2024.
- [28] Hao Xu, Peize Liu, Xinyi Chen, and Shaojie Shen. d^2 slam: Decentralized and distributed collaborative visual-inertial slam system for aerial swarm. *IEEE Transactions on Robotics*, 2024.
- [29] Hao Xu, Yichen Zhang, Boyu Zhou, Luqi Wang, Xinjie Yao, Guotao Meng, and Shaojie Shen. Omni-swarm: A decentralized omnidirectional visual-inertial-uwB state estimation system for aerial swarms. *Ieee transactions on robotics*, 38(6):3374–3394, 2022.
- [30] Zhiren Xunt, Jian Huang, Zhehan Li, Zhenjun Ying, Yingjian Wang, Chao Xu, Fei Gao, and Yanjun Cao. Crepes: Cooperative relative pose estimation system. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5274–5281. IEEE, 2023.
- [31] Heng Yang, Jingnan Shi, and Luca Carlone. Teaser: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, 37(2):314–333, 2021.
- [32] Peng Yin, Lingyun Xu, Ji Zhang, Howie Choset, and Sebastian Scherer. i3dloc: Image-to-range cross-domain localization robust to inconsistent environmental conditions. In *17th Robotics: Science and Systems, RSS 2021*, 2021.
- [33] Chunran Zheng, Wei Xu, Zuhao Zou, Tong Hua, Chongjian Yuan, Dongjiao He, Bingyang Zhou, Zheng Liu, Jiarong Lin, Fangcheng Zhu, et al. Fast-livo2: Fast, direct lidar-inertial-visual odometry. *IEEE Transactions on Robotics*, 2024.
- [34] Fangcheng Zhu, Yunfan Ren, Fanze Kong, Huajie Wu, Siqi Liang, Nan Chen, Wei Xu, and Fu Zhang. Swarm-liv: Decentralized swarm lidar-inertial odometry. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3254–3260. IEEE, 2023.