

CU-Multi: A Dataset for Multi-Robot Collaborative Perception

Doncey Albin¹, Daniel McGann², Miles Mena¹, Annika Thomas³, Harel Biggie⁴, Xuefei Sun¹
Steve McGuire⁵, Jonathan P. How³, and Christoffer Heckman¹

Abstract—A central challenge for multi-robot systems is fusing independently gathered perception data into a unified representation. Despite progress in Collaborative SLAM (C-SLAM), benchmarking remains hindered by the scarcity of dedicated multi-robot datasets. Many evaluations instead partition single-robot trajectories, a practice that may only partially reflect true multi-robot operations and, more critically, lacks standardization, leading to results that are difficult to interpret or compare across studies. While several multi-robot datasets have recently been introduced, they mostly contain short trajectories with limited inter-robot overlap and sparse intra-robot loop closures. To overcome these limitations, we introduce CU-Multi, a dataset collected over multiple days at two large outdoor sites on the University of Colorado Boulder campus. CU-Multi comprises four synchronized runs with aligned start times and controlled trajectory overlap, replicating the distinct perspectives of a robot team. It includes RGB-D sensing, RTK GPS, semantic LiDAR, and refined ground-truth odometry. By combining overlap variation with dense semantic annotations, CU-Multi provides a strong foundation for reproducible evaluation in multi-robot collaborative perception tasks. The dataset, support code, and updates are available at <https://arpg.github.io/cumulti>.

I. INTRODUCTION

Multi-robot systems significantly enhance capabilities across diverse domains, particularly in large-scale environments, by accelerating exploration through distributed sensing and collaborative decision-making [7]. A central challenge in realizing these advantages lies in fusing perception data collected independently by multiple robots into a unified global representation, complicated by spatial and temporal misalignment. This challenge is further compounded in many practical scenarios where external positioning methods (e.g. GPS, motion capture) are impractical, unreliable, or hazardous. Multi-robot Collaborative SLAM (C-SLAM) algorithms are typically first developed and validated offline using datasets before doing field tests in real-world conditions, where system-level issues often arise. Consequently, the availability of realistic, well-structured multi-robot datasets is essential for supporting reproducible

¹Authors are with the Autonomous Robotics and Perception Group at the University of Colorado Boulder, Boulder, CO 80309, USA

²Daniel McGann with the Robot Perception Lab at Carnegie Mellon University, Pittsburgh, PA 15213, USA

³Authors are with the Aerospace Controls Laboratory at Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁴Harel Biggie with the Computer Science and Artificial Intelligence Laboratory at Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁵Steve McGuire with the Human-Aware Robotic Exploration Lab at University of California Santa Cruz, Santa Cruz, CA 95064, USA

This work was supported in part by USDA-NIFA award 2021-67021-33450 and by the National Science Foundation (NSF) award #2339328.

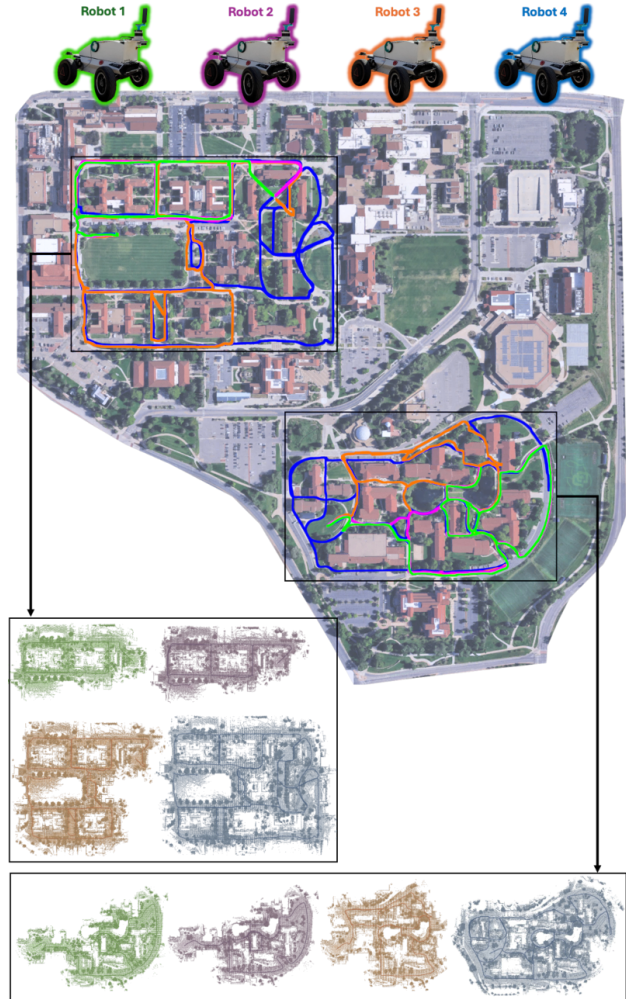


Fig. 1. Overhead view of all paths overlaid on map from the Main Campus (top) and Kitteridge Loop (bottom) environments in the CU-Multi Dataset.

research and bridging the gap between offline development and real-world deployment.

Several multi-robot datasets have recently been introduced for C-SLAM verification [1], [3], [6]. These datasets include synchronized multi-agent sensor data collected across both indoor and/or outdoor environments. While these datasets represent strong starting points, there remains a need for longer-trajectory multi-robot datasets with explicitly varied overlap (see Table I). Despite recent advancements in multi-robot datasets, there remains a prevalent practice of artificially segmenting a single trajectory into multiple parts to simulate a multi-robot scenario for verification [8]. Single-

TABLE I
COMPARISON OF FIELD-COLLECTED MULTI-ROBOT DATASETS. NOTE: ONLY DATASETS WITH LiDAR HAVE BEEN LISTED.

Dataset	# Robots	Environments	Multi-Session	RGB-D	GPS	IMU	Annotated LiDAR	Longest Trajectory (km)
GRACO [1]	1 GW, 2 A	8 Out	✗	✓	✓	✓	✗	0.82
Kimera-Multi [2]	8 GW	1 Out, 1 In, 1 Hybrid	✗	✓	✗	✓	✗	1.40
S3E [3]	3 GW	13 Out, 5 In	✗	✓	✓	✓	✗	1.95
DiTer++ [4]	2 GL	3 Out	✓	✓	✗	✓	✗	N/A
Lamp 2.0 [5]	4 GW, 3 GL	1 Out, 3 Sub	✗	✓	✗	✓	✗	2.20
CoPeD [6]	3 GW, 2 A	3 In, 3 Out	✗	✓	✓	✓	✗	0.36
CU-Multi (ours)	4 GW	2 Out	✓	✓	✓	✓	✓	4.01

GW, GL, and A denote ground-wheeled, ground-legged, and aerial robots, respectively. In, Out, and Hybrid represent indoor, outdoor, and a mix of both, respectively.

robot SLAM datasets still offer a few benefits over many of the existing multi-robot datasets, such as large trajectories with multiple loop closures, as well as the availability of both camera and LiDAR semantics [8], [9]. However, without careful consideration, arbitrary segmentation may not accurately represent realistic observational overlap typically encountered in multi-robot operations [10].

In this paper, we introduce CU-Multi, a dataset designed to support the evaluation of methods relevant to multi-robot perception. These methods include C-SLAM, LiDAR and visual data association, and multi-session LiDAR-based place recognition. CU-Multi offers the following key features:

- 1) *A multi-robot dataset* consisting of two large-scale environments, each with four robots. Our dataset contains a total of eight diverse trajectories, spanning a combined length of 16.7 km across the CU Boulder campus (Figure 1).
- 2) *Systematically varied trajectory overlaps and a rendezvous-based trajectory design*, enabling evaluation of multi-robot perception tasks under different levels of observational redundancy and variation.
- 3) *Refined ground truth poses* using RTK GPS, a digital elevation model, LiDAR-inertial SLAM, and a highly accurate scan-matching algorithm, calculated at each LiDAR timestamp.
- 4) *Semantic labels* for all LiDAR scans, dataset tools for interacting with the data, as well as initial benchmarks on a well-known C-SLAM and LiDAR place recognition method to demonstrate dataset usability and utility of the provided tools.

II. RELATED WORK

Although single-robot SLAM has long benefited from curated datasets, extending SLAM to multi-robot systems introduces challenges that demand more representative data of a true multi-robot system. A common practice is to simulate multi-robot scenarios by splitting single-robot trajectories, but this raises concerns about realism and consistency. This section surveys single-robot datasets, the trajectory-splitting strategies applied to them, and existing multi-robot datasets, motivating the need for our proposed contribution.

A. Single-Robot SLAM Datasets

One of the most widely used datasets for evaluating single-robot SLAM and global localization pipelines is the KITTI Odometry Benchmark [11], a large-scale computer vision dataset collected in Karlsruhe, Germany using a sensor-equipped automobile. It features stereo imagery, LiDAR, and GPS/IMU data across 22 driving sequences in urban and rural environments. In 2019, SemanticKITTI [12] expanded the original KITTI dataset to include dense point-wise annotations for all LiDAR scans, enabling tasks such as semantic segmentation and scene completion. KITTI-360 [13], released in 2022, was also collected in Karlsruhe and contains broader city coverage, improved GPS accuracy, and consistent 2D/3D semantic instance annotations across panoramic imagery and LiDAR scans. Although originally intended for single-robot perception tasks, these datasets are commonly repurposed for verifying and benchmarking multi-robot C-SLAM pipelines through segmenting a single trajectory into multiple parts to simulate independent robots.

B. Trajectory Splitting on Single-Robot Datasets

Early work by Cieslewski and Scaramuzza demonstrated the approach of using a single-robot dataset for C-SLAM [14], [15], and subsequent methods including DOOR-SLAM [16], DiSCO-SLAM [17], and RDC-SLAM [18] adopted similar practices. More recent systems, such as DCL-SLAM [8], Wi-Closure [19], and SWARM-SLAM [9], continue to rely on partitioning KITTI or KITTI-360 sequences, often supplemented with small multi-robot datasets like GRACO [1] or S3E [3].

While splitting provides convenient large-scale data, its limitations are well documented. Lajoie et al. cautioned that partitioning trajectories at overlapping regions creates unrealistic assumptions, as identical viewpoints and lighting are unlikely to occur simultaneously in true multi-robot operations [10]. Despite this, most works do not specify their partitioning strategies, and no standardized protocol exists. As seen in methods such as FRAME [20], SideSLAM [21], and Cao et al. [22], overlapping splits can even reuse the same observations, undermining realism. This persistent reliance highlights a key gap: the lack of standardized

datasets that capture the diversity and independence of multi-robot observations.

C. Multi-Robot Datasets

One of the earliest multi-robot datasets is the 2011 UTIAS dataset [23], which includes nine indoor environments, each featuring five wheeled robots equipped with monocular cameras. In 2020, two additional datasets were introduced: FordAV [24] and AirMuseum [25]. FordAV provides multi-seasonal data collected by a fleet of Ford Fusion vehicles outfitted with four LiDARs, GPS, and six cameras. However, the vehicle trajectories exhibit near-complete overlap across all sequences, requiring users to manually segment them to simulate limited overlap scenarios. AirMuseum captures warehouse-scale data using a heterogeneous team of two ground robots and one aerial platform, each equipped with a stereo camera rig. While both AirMuseum and UTIAS offer multi-robot observations, their indoor collection settings and lack of LiDAR data limit their scalability and applicability to modern perception pipelines.

More recent datasets have broadened the scope of multi-robot data collection. GRACO [1] introduced a heterogeneous two-robot team consisting of a ground and aerial robot; however, the dataset is constrained by a low-resolution 16-beam LiDAR and short trajectories. In follow-up work to Kimera-Multi [26], Kim et al. released a dataset comprising three sequences of eight ground robots operating in indoor, outdoor, and hybrid environments [2]. The CoPeD dataset goes further by addressing specific limitations of prior datasets in terms of sensor heterogeneity and environmental diversity [6]. Kimera-Multi and CoPeD address valuable aspects needed for a multi-robot dataset, but are still limited to small trajectories and minimal observational overlap.

The need for datasets that explicitly explore varying degrees of shared observations has also gained attention. S3E presents four distinct trajectory paradigms to study different levels of inter-robot trajectory overlap using a team of three robots [3]. In their work, they intended to reduce the amount of inter-robot trajectory overlap to simulate scenarios that are not favorable. While this approach is a valuable step toward evaluating C-SLAM methods under varied trajectory overlap conditions, the trajectories remain limited in scale and contain relatively few intra-robot loop closures.

Despite recent progress in multi-robot datasets, many works have yet to adopt them for evaluation. Relying on idealized, single-trajectory conditions can obscure failure modes that emerge when robots operate independently, with varied trajectory overlap, differing viewpoints, and asynchronous data capture. As research in multi-robot perception continues to mature, there is a growing need for datasets that provide precise ground truth across multiple independently moving platforms to rigorously evaluate C-SLAM methods and inter-robot loop closure detection. Our dataset addresses this gap by offering multi-robot trajectories with explicitly controlled overlap levels and annotated LiDAR scans, enabling more realistic and robust evaluation of multi-robot perception pipelines.




Fig. 2. Hardware and sensor specifications used on platform.

III. THE CU-MULTI DATASET

The CU-Multi dataset was collected using a single robotic platform over multiple sessions. This section describes the platform and sensors used for data collection, our novel approach to ensuring accurate ground-truth localization across multiple runs within each environment, our two-stage method for LiDAR annotation, the large-scale outdoor environments where data was gathered, and how the dataset is structured.

A. Platform and Sensors

The CU-Multi dataset was collected using the AgileX Hunter SE platform modified with a custom-designed electronics housing (see Figure 2). Inside the housing, the system includes an Intel NUC i7 with 16 GB of RAM, a power distribution board, and a 217 Wh battery. Externally, a mounting plate on the rear of the Hunter SE accommodates two u-blox ANN-MB-00 GNSS antennas, while a front-mounted plate holds a 64-beam Ouster LiDAR sensor, a




Fig. 3. During data collection, we ensured each session started from the same initial starting point for each environment through securing our platform onto four wheel chocks that were staked into grass (left). During data playback, we can see the results of our initial alignment through tightly overlapping lidar scans (top right). Lastly, we ended each session at a common location, each run ending within 5m of one another (bottom right).

RealSense D455 RGB-D camera, a Lord MicroStrain G7 IMU, and a Lord RTK cellular modem with an external antenna. To ensure precise odometry estimation, LiDAR timestamps are hardware synchronized with the onboard computer’s timeclock using Precision Time Protocol (PTP).

B. Dataset Collection and Environments

The CU-Multi dataset was collected across two large outdoor environments on the University of Colorado Boulder campus. The first environment, *main_campus*, covers the central academic area, with approximately 7.4 km of traversed path. The second environment, *kittredge_loop*, encompasses a large region south of the main campus, offering more open space and varied terrain for multi-robot exploration.

Each of the two environments consists of four trajectories with varying levels of overlap. In our dataset, *robot1* and *robot2* share significant path overlap but capture the scene from different viewpoints, enabling users to test algorithms for robust multi-session data association. Meanwhile, *robot3* and *robot4* follow paths with less overlap, providing opportunities to evaluate performance under sparse or partially overlapping observations. The trajectory of *robot3* extends the trajectory of the first two robots, introducing additional observations. Finally, *robot4* encompasses all previous trajectories, covering the most extensive area. Formally, the relationship between these trajectories can be expressed as:

$$T_1 \approx T_2, \quad (T_1 \cup T_2) \subseteq T_3, \quad (T_1 \cup T_2 \cup T_3) \subseteq T_4 \quad (1)$$

where T_i represents the trajectory of *robot_i*.

To maximize viewpoint diversity, we introduced variations in observational perspectives while traversing overlapping regions. The structured overlap across trajectories enables users to select data subsets based on their desired level of spatial redundancy and cross-view consistency. All trajectories end within a distance of 5 meters of each other at a common location in each environment, representing a multi-robot rendezvous scenario (Figure 3).

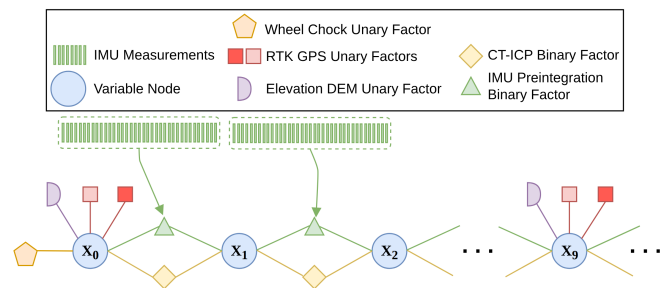


Fig. 4. The factor-graph used to generate ground truth for all trajectories. The LiDAR and IMU factors interpolate between the poses that are well constrained by external reference data (i.e. GPS, Elevation, Chock Prior).




Fig. 5. The GPS uncertainty in both data collection environments reported as $\|\Sigma\|_F$. Covariances (Σ) are reported by the GNSS and are spatially averaged across all trajectories using a 4m resolution voxel-grid.

C. Ground Truth

To support quantitative evaluation, we provide globally aligned ground truth for all trajectories. We initialized each trajectory from a fixed point by securing the Hunter SE platform onto wheel chocks staked into the grass (Figure 3), after which the robot remained stationary for ~ 60 s to acquire an accurate RTK GPS fix. However, due to the density of buildings in traversed sections of the environments, the RTK GPS does lose some accuracy, particularly in elevation (see Figure 5), and at 2Hz the GPS data alone is too sparse to provide useful ground truth.

To address these limitations, we formulate a factor-graph incorporating GPS measurements from both GNSS receivers,

20Hz LiDAR odometry factors from LIO-SAM [27] and CT-ICP [28], and IMU preintegration factors to interpolate between GPS fixes. Elevation priors from a digital elevation model (DEM)¹ mitigate large elevation ambiguity in dense areas, and a unary prior at the wheel chock location constrains each trajectory’s initial pose. The resulting graph is optimized using `gtsam` [30], yielding a high-quality pose estimate at every LiDAR scan timestamp. We qualitatively validate these estimates by overlaying the reconstructed maps onto an independently acquired digital surface model (DSM)² collected over the same area, observing strong spatial agreement across both environments (Fig. 6).




Fig. 6. Maps constructed from the estimated ground truth poses overlaid on an external DSM for qualitative validation. Maps are colored per robot. Top: Kittredge Loop (perspective). Bottom: Main Campus (left) and Kittredge Loop (right) in overhead view.

¹We use the 1m-resolution DEM of Boulder, CO provided by the USGS [29].

²Derived from the DRCOG 2020 QL1 LiDAR collection [31].

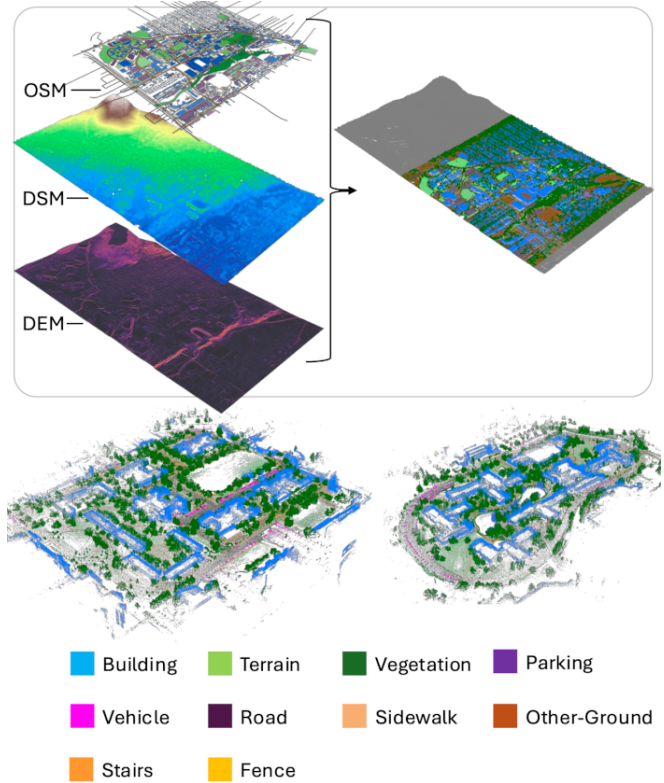


Fig. 7. Illustration of the automated semantic labeling pipeline. External geospatial data sources, including OpenStreetMap (OSM) features, a digital surface model (DSM), and a digital elevation model (DEM), are fused to construct a global semantic map (top). This semantic prior is then used to generate semantically annotated LiDAR scans, which are aggregated into semantic environment maps for the Main Campus (bottom left) and Kittredge Loop (bottom right) sequences.

D. LiDAR Annotation

None of the multi-robot datasets surveyed in this work provide ground-truth semantic annotations for LiDAR data, although CoPeD [6] does include semantic labels for images. Generating LiDAR semantic annotations at scale is expensive and remains a major obstacle to building such datasets. To address this gap, we develop an automated labeling pipeline that combines accurate ground-truth odometry with open geospatial data, including the DEM and DSM described in Section III-C, as well as OpenStreetMap (OSM) features (Fig. 7).

First, we use the DEM to identify which DSM points belong to the ground surface, complementing the limited class information available in the DSM itself. We then use semantic OSM features to label the DSM, producing a semantic DSM whose annotations are constrained by whether points correspond to ground-level regions or above-ground structures. Finally, we transfer these labels to each LiDAR scan by assigning to every LiDAR point the label of its nearest neighbor in the semantic DSM. This procedure is enabled by the globally aligned, accurate ground-truth odometry, which allows reliable registration between the sensor data and the geospatial priors.

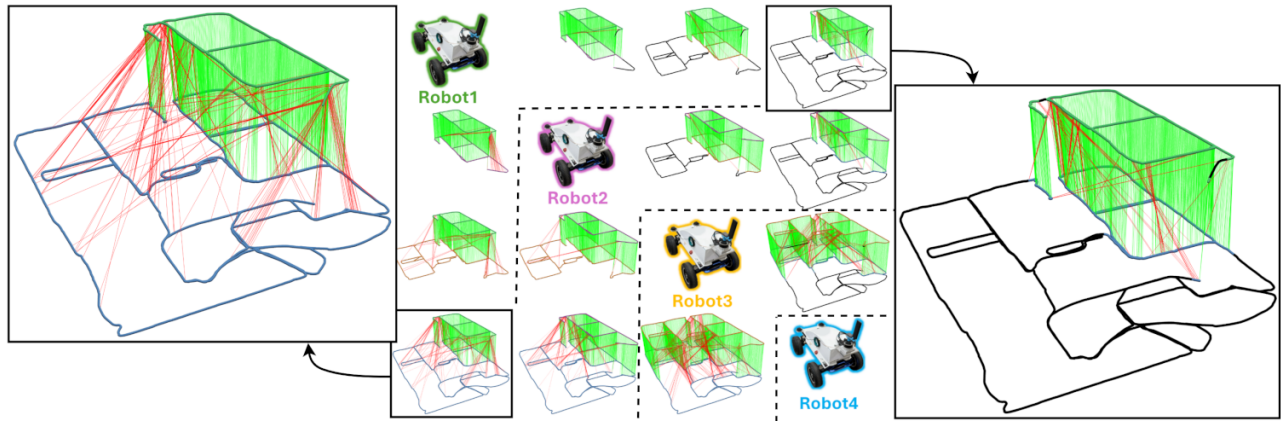


Fig. 8. Example demonstrating utilization of LiDAR place recognition using ScanContext on data collected in Main Campus environment. All examples above the main diagonal compare multi-session place recognition and all examples below are multi-robot comparisons. Red lines show false matching loops and green lines show correct matching loops. We determine a loop-closure is correct if the positions of the scans are within 10m.

E. Dataset Structure and Format

The dataset is structured hierarchically, with environments and a calibration directory at the top level. The calibration directory (*calib*) contains the relative sensor positions on the platform, camera intrinsics, a mesh and URDF of the system. Within each environment directory, there are four subdirectories corresponding to individual robots, labeled *robot1* through *robot4*.

Each robot-specific directory follows a consistent structure with subdirectories containing ROS2 *.db3* bags for the onboard RGB-D camera, LiDAR, IMU, GPS, and ground truth poses. The ground truth poses in the ROS2 *.db3* bags are centered about the starting position, therefore we also include a *.csv* in each robot subdirectory with the ground truth poses in UTM coordinates with each line formatted as $[timestamp, x, y, z, qx, qy, qz, qw]$. Lastly, on our project website, we illustrate the topics available to users in each of the ROS2 bags, which sensor (if applicable) collected the topic data, the rate of that topic, and the ROS2 message type in which the corresponding data are stored.

F. Dataset Tools

To facilitate adoption and reproducible experimentation, CU-Multi is released with a suite of tools. These include Python scripts for converting the raw data into common formats such as ROS1 bag files and ROS2 *.mcap* files. We also provide an interactive trajectory clipper tool that allows users to select, crop, and replay robot trajectories. This enables controlled-overlap scenarios beyond those observed during data collection, while still using runs that were independently captured with varied viewpoints. Such functionality is particularly useful for evaluating C-SLAM methods that employ distance-based communication constraints, where playback overlap can ensure that robot observations occur within a desired proximity. Together, these tools streamline integration with existing multi-robot perception pipelines, enabling rapid and flexible evaluation. Documentation and usage examples are provided on the project website.

IV. USE CASE DEMONSTRATIONS

To demonstrate CU-Multi's suitability for multi-robot C-SLAM and data association, we present two illustrative evaluations. For this purpose, we select two lightweight baselines: the LiDAR global descriptor ScanContext [32] and the distributed C-SLAM system DiSCo-SLAM. Since DiSCo-SLAM employs ScanContext for inter-robot loop closure detection, we apply identical ScanContext parameters in both evaluations to ensure consistency. These parameters, along with step-by-step instructions, are released with our code to support reproducibility and facilitate drop-in integration of future methods.

A. LiDAR Place Recognition

ScanContext constructs a rotation-tolerant global descriptor for each LiDAR scan to enable fast retrieval. We evaluate precision-recall in both multi-session and multi-robot settings for each CU-Multi environment. In the multi-session setting, we compare descriptors only between pose pairs whose ground-truth separation is at most 10m across sessions. In the multi-robot setting, we perform all-pairs retrieval: all scans from one robot are compared against all scans from another robot (Fig. 8). We provide PR curves for both the multi-robot and multi-session place recognition experiments (Fig. 10).

B. LiDAR-Based C-SLAM

DiSCo-SLAM exchanges compact ScanContext descriptors and performs a two-stage global-local pose-graph optimization for distributed multi-robot SLAM. The reference implementation supports up to three robots, so we extended it to handle all four robots available in CU-Multi. We tested DiSCo-SLAM four times in the Main Campus environment and five times in the Kittredge Loop environment. Of these trials, one run on Main Campus and two runs on Kittredge Loop diverged due to erroneous loop closures. To ensure that inter-robot loop closures did not occur immediately, playback was offset by 200 seconds from the initial start time. We report the averages for Absolute Trajectory Error (ATE) and Relative Pose Error (RPE) on all runs that remained stable




Fig. 9. *DiSCo-SLAM* on *CU-Multi*. **Top row**: Main Campus environment. **Bottom row**: Kittredge Loop environment. Each row illustrates four stages of alignment from left to right: (1) initial unaligned robot trajectories, (2) *robot2* aligned to *robot1*, (3) *robot3* subsequently aligned with *robot1* and *robot2*, and (4) *robot4* aligned with the remaining robots, resulting in all trajectories progressively unified into a common frame.

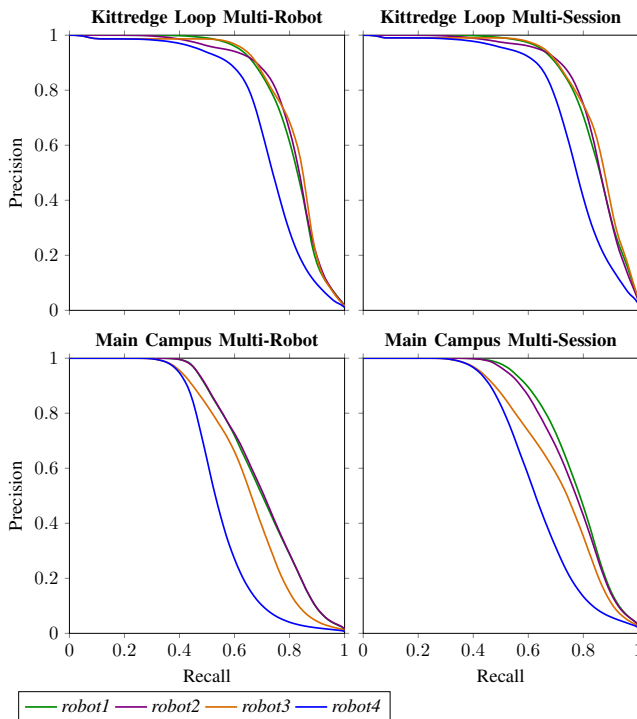


Fig. 10. Precision-Recall curves for the ScanContext descriptor for LiDAR-based place recognition for both environments in our dataset using a multi-session test and a multi-robot test.

and did not fail due to incorrect loop closure detections (Table II). Figure 9 illustrates an example of progressive inter-robot alignment produced by *DiSCo-SLAM* across both environments.

C. Results

From our experiments, we found that on high-overlap pairs (e.g., *robot1* and *robot2*), *DiSCo-SLAM* rapidly establishes inter-robot loop closures leading to better overall

TABLE II
DiSCo-SLAM RESULTS ON THE CU-MULTI DATASET

Environment		robot1	robot2	robot3	robot4
Main Campus	<i>ATE</i> (m)	5.58	8.52	14.30	22.85
	<i>RPE</i> (m)	0.11	0.11	0.49	0.31
Kittredge Loop	<i>ATE</i> (m)	7.68	5.01	14.42	18.85
	<i>RPE</i> (m)	0.08	0.16	0.35	0.45

ATE and ScanContext yields stronger precision-recall. As overlap decreases (*robot1* and *robot4*), PR degrades and time-to-first closure increases, quantifying the practical effect of viewpoint diversity and partial overlap that *CU-Multi* was designed to expose. These results, together with our globally aligned ground truth and semantics, showcase *CU-Multi* as the right substrate for realistic evaluation without resorting to single-trajectory splitting.

V. CONCLUSION AND FUTURE WORK

We introduced *CU-Multi*, a multi-robot dataset that addresses persistent challenges in evaluating collaborative perception and SLAM. By combining diverse trajectories with controlled overlap, synchronized multi-modal sensing, semantic LiDAR annotations, and refined ground-truth odometry across two large-scale outdoor environments, *CU-Multi* enables realistic and reproducible experimentation beyond single-robot trajectory partitioning. Baseline evaluations confirm that varied overlap levels expose meaningful performance differences in both place recognition and C-SLAM. The public release of the dataset, support code, and documentation is intended to accelerate progress in collaborative autonomy. In future work, we plan to extend *CU-Multi* to include heterogeneous platforms and indoor environments.

REFERENCES

- [1] Y. Zhu, Y. Kong, Y. Jie, S. Xu, and H. Cheng, "GRACO: A multimodal dataset for ground and aerial cooperative localization and mapping," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 966–973, 2023.
- [2] Y. Tian, Y. Chang, L. Quang, A. Schang, C. Nieto-Granda, J. How, and L. Carlone, "Resilient and distributed multi-robot visual SLAM: Datasets, experiments, and lessons learned," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2023.
- [3] D. Feng, Y. Qi, S. Zhong, Z. Chen, Q. Chen, H. Chen, J. Wu, and J. Ma, "S3E: A Multi-Robot Multimodal Dataset for Collaborative SLAM," *IEEE Robotics and Automation Letters*, 2024.
- [4] J. Kim, H. Kim, S. Jeong, Y. Shin, and Y. Cho, "DiTer++: Diverse Terrain and Multi-modal Dataset for Multi-Robot SLAM in Multi-session Environments," *arXiv preprint arXiv:2412.05839*, 2024.
- [5] Y. Chang, K. Ebadi, C. E. Denniston, M. F. Ginting, A. Rosinol, A. Reinke, M. Palieri, J. Shi, A. Chatterjee, B. Morrell *et al.*, "LAMP 2.0: A robust multi-robot SLAM system for operation in challenging large-scale underground environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9175–9182, 2022.
- [6] Y. Zhou, L. Quang, C. Nieto-Granda, and G. Loianno, "CoPeD-Advancing Multi-Robot Collaborative Perception: A Comprehensive Dataset in Real-World Environments," *IEEE Robotics and Automation Letters*, 2024.
- [7] S. Kim, M. Corah, J. Keller, G. Best, and S. Scherer, "Multi-robot multi-room exploration with geometric cue extraction and circular decomposition," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1190–1197, 2023.
- [8] S. Zhong, Y. Qi, Z. Chen, J. Wu, H. Chen, and M. Liu, "DCL-SLAM: A distributed collaborative lidar SLAM framework for a robotic swarm," *IEEE sensors journal*, vol. 24, no. 4, pp. 4786–4797, 2023.
- [9] P.-Y. Lajoie and G. Beltrame, "Swarm-slam: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 475–482, 2023.
- [10] P.-Y. Lajoie, B. Ramtoul, F. Wu, and G. Beltrame, "Towards collaborative simultaneous localization and mapping: a survey of the current research landscape," *Field Robotics*, vol. 2, pp. 971–1000, 2022.
- [11] A. Geiger, P. Lenz, C. Still, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [12] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "Semantickitti: A dataset for semantic scene understanding of lidar sequences," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9297–9307.
- [13] Y. Liao, J. Xie, and A. Geiger, "Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2D and 3D," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 3292–3310, 2022.
- [14] T. Cieslewski and D. Scaramuzza, "Efficient decentralized visual place recognition using a distributed inverted index," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 640–647, 2017.
- [15] T. Cieslewski, S. Choudhary, and D. Scaramuzza, "Data-efficient decentralized visual SLAM," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 2466–2473.
- [16] P.-Y. Lajoie, B. Ramtoul, Y. Chang, L. Carlone, and G. Beltrame, "DOOR-SLAM: Distributed, online, and outlier resilient SLAM for robotic teams," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1656–1663, 2020.
- [17] Y. Huang, T. Shan, F. Chen, and B. Englot, "DiSCo-SLAM: Distributed scan context-enabled multi-robot lidar SLAM with two-stage global-local graph optimization," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1150–1157, 2021.
- [18] Y. Xie, Y. Zhang, L. Chen, H. Cheng, W. Tu, D. Cao, and Q. Li, "RDC-SLAM: A real-time distributed cooperative SLAM system based on 3D lidar," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 9, pp. 14721–14730, 2021.
- [19] W. Wang, A. Kemmeren, D. Son, J. Alonso-Mora, and S. Gil, "Wireless closure: Reliable and efficient search of inter-robot loop closures using wireless sensing," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 2069–2075.
- [20] N. Stathouloupoulos, B. Lindqvist, A. Koval, A.-A. Agha-Mohammadi, and G. Nikolakopoulos, "FRAME: a modular framework for autonomous map-merging: Advancements in the field," *IEEE Transactions on Field Robotics*, 2024.
- [21] X. Liu, J. Lei, A. Prabhu, Y. Tao, I. Spasojevic, P. Chaudhari, N. Atanasov, and V. Kumar, "Slideslam: Sparse, lightweight, decentralized metric-semantic SLAM for multi-robot navigation," *arXiv preprint arXiv:2406.17249*, 2024.
- [22] H. Cao, S. Shreedharan, and N. Atanasov, "Multi-robot object SLAM using distributed variational inference," *IEEE Robotics and Automation Letters*, 2024.
- [23] K. Y. Leung, Y. Halpern, T. D. Barfoot, and H. H. Liu, "The UTIAS multi-robot cooperative localization and mapping dataset," *The International Journal of Robotics Research*, vol. 30, no. 8, pp. 969–974, 2011.
- [24] S. Agarwal, A. Vora, G. Pandey, W. Williams, H. Kourous, and J. McBride, "Ford multi-av seasonal dataset," *The International Journal of Robotics Research*, vol. 39, no. 12, pp. 1367–1376, 2020.
- [25] R. Dubois, A. Eudes, and V. Frémont, "AirMuseum: a heterogeneous multi-robot dataset for stereo-visual and inertial simultaneous localization and mapping," in *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2020, pp. 166–172.
- [26] Y. Tian, Y. Chang, F. H. Arias, C. Nieto-Granda, J. P. How, and L. Carlone, "Kimera-multi: Robust, distributed, dense metric-semantic SLAM for multi-robot systems," *IEEE Transactions on Robotics*, vol. 38, no. 4, 2022.
- [27] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and R. Daniela, "LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5135–5142.
- [28] P. Dellenbach, J.-E. Deschaud, B. Jacquet, and F. Goulette, "CT-ICP: real-time elastic lidar odometry with loop closure," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 5580–5586.
- [29] U.S. Geological Survey, "USGS 1 meter 13 x47y443 CO_DRCOG_2020_B20," 2022, www.sciencebase.gov/catalog/item/620de541d34e6c7e83baa08a.
- [30] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," Georgia Institute of Technology, Technical Report, 2012.
- [31] DRCOG Regional Data Catalog, "LiDAR QL1 Index in CO SP Central 2020," <https://data.drcog.org/dataset/lidar-ql1-index-in-co-sp-central-2020>, 2020, licensed under CC BY 3.0. Accessed: 2025.
- [32] G. Kim and A. Kim, "Scan context: Egocentric spatial descriptor for place recognition within 3D point cloud map," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4802–4809.