

Hilti-Oxford Dataset: A Millimeter-Accurate Benchmark for Simultaneous Localization and Mapping

Lintong Zhang¹, Michael Helmberger², Lanke Frank Tarimo Fu¹, David Wisht¹, Marco Camurri¹, Davide Scaramuzza³, Maurice Fallon¹

Abstract—Simultaneous Localization and Mapping (SLAM) is being deployed in real-world applications, however many state-of-the-art solutions still struggle in many common scenarios. A key necessity in progressing SLAM research is the availability of high-quality datasets and fair and transparent benchmarking. To this end, we have created the Hilti-Oxford Dataset, to push state-of-the-art SLAM systems to their limits. The dataset has a variety of challenges ranging from sparse and regular construction sites to a 17th century neoclassical building with fine details and curved surfaces. To encourage multi-modal SLAM approaches, we designed a data collection platform featuring a lidar, five cameras, and an IMU (Inertial Measurement Unit). With the goal of benchmarking SLAM algorithms for tasks where accuracy and robustness are paramount, we implemented a novel ground truth collection method that enables our dataset to accurately measure SLAM pose errors with millimeter accuracy. To further ensure accuracy, the extrinsics of our platform were verified with a micrometer-accurate scanner, and temporal calibration was managed online using hardware time synchronization. The multi-modality and diversity of our dataset attracted a large field of academic and industrial researchers to enter the second edition of the Hilti SLAM challenge, which concluded in June 2022. The results of the challenge show that while the top three teams could achieve an accuracy of 2 cm or better for some sequences, the performance dropped off in more difficult sequences.

Index Terms—Data Sets for SLAM; SLAM; Mapping;

I. INTRODUCTION

SLAM research has made impressive progress, allowing the transition from lab demonstrations to real-world deployment. Open-source datasets play a key role in this transition, as researchers can progressively improve and compare different SLAM solutions. The TUM [1], EuRoC [2], and KITTI [3] datasets have been a pillar in the robotics community, and their leaderboards are still motivating new and improved algorithms. However, as SLAM algorithms improve to enable real-life applications, so should their benchmarks. We see the need to have more challenging datasets to differentiate top SLAM approaches. We also believe that these datasets should use the latest sensors to benchmark multi-modal SLAM accuracy and robustness under a variety of challenges.

The key motivation for this work is to create a high-quality dataset with a variety of challenging sequences that can propel SLAM-related research. As shown in several works [4], [5], [6], it is beneficial to fuse multiple sensors to improve accuracy

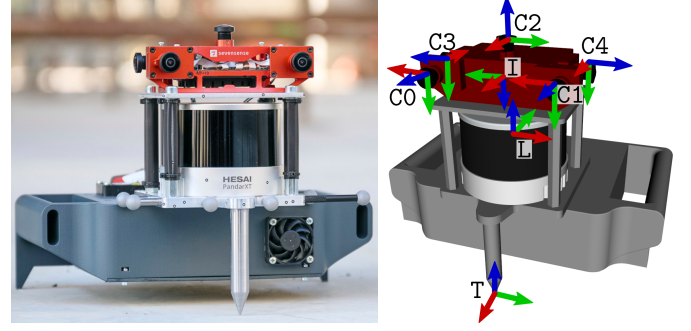


Fig. 1: The handheld device, called Phasma, is composed of five cameras, an IMU sensor, and a 32 beam lidar.

and robustness. Hence, we present a SLAM dataset combining vision, lidar, and inertial sensing.

Additionally, accurate ground truth plays a vital role in evaluation, as many lidar-based algorithms are approaching centimeter-level accuracy. Thus, we propose a new approach which can provide millimeter accuracy by using state-of-art surveying equipment. By leveraging a high precision long-range lidar, multiple global shutter cameras, and a synchronized IMU, we provide a comprehensive dataset where SLAM algorithms must perform accurately and robustly in order to be used in real-world applications.

The need for and interest in a high-quality dataset is evidenced by the Hilti SLAM Challenge 2022¹, which received 42 submissions from both industry and academic research groups. In summary, the Hilti-Oxford dataset² offers the following contributions:

- Challenging and degenerate scenarios, such as dark corners, narrow stairs, long corridors, and a few dynamic objects, with sequences specifically designed to break existing SLAM algorithms;
- A data collection platform with modern sensors, including an accurate long-range lidar (up to 120 m), five fisheye cameras operating at 40 Hz, and inertial sensors. The sensors are mounted on a high-precision machined chassis, with extrinsics verified by a micrometer accurate scanner. All signals have been hardware synchronized;
- A novel sparse ground truth collection method based on a survey-grade scanner and reference targets, which achieves millimeter precision;
- Insights and discussion of the merits of each system and sensor modality based on the high number of submissions.

¹Oxford Robotics Institute, Department of Engineering Science, University of Oxford, UK. lintong@robots.ox.ac.uk

²Hilti AG, Schaan, Liechtenstein

³Robotics and Perception Group, Department of Informatics, University of Zurich, Switzerland

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

¹Challenge Video: <https://www.youtube.com/watch?v=-LMq3zU47Pw>

²<https://hilti-challenge.com/dataset-2022.html>

II. RELATED WORK

Existing benchmark datasets for SLAM can be categorized by their different operation domains. Depending on the domain, different sensory data and varying degrees of ground truth accuracy are provided.

In the visual-inertial odometry domain, EuRoC [2] and TUM VI [1], which provide camera and IMU data, have been extensively used by the research community. EuRoC recorded hardware time-synchronized stereo camera images and IMU measurements from a micro aerial vehicle equipped with a Skybotix stereo VI sensor. While the 11 EuRoC sequences are accompanied by millimeter accurate ground truth poses, their trajectories only covered the indoor bounds of a motion capture system. The TUM VI dataset uses a hand-held data acquisition device which features a global shutter stereo camera and IMU, which were also hardware time-synchronized. Unlike EuRoC, the 28 sequences in TUM VI include segments that extend outdoors, providing a diverse set of scenarios to benchmark visual-inertial SLAM. The segments in TUM VI that extend outdoors start and end in an indoor motion capture environment, providing ground truth poses in these indoor segments.

Datasets such as KITTI [3], WoodScape [7], and UMich [8] are specialized to the autonomous driving domain, and in addition to IMU and camera images, these datasets also provide lidar data. WoodScape and UMich are both equipped with 360° cameras while KITTI used a linear array of stereo cameras, two colored and two grayscale. With a Segway scooter as a platform, the UMich dataset covers both indoor and outdoor scenarios. KITTI and WoodScape, on the other hand, used large vehicles as their platform, so were not able to record indoor environments. All these datasets (KITTI, WoodScape, and UMich) provided ground truth trajectories using GPS/GNSS-INS measurements which are only accurate to several centimeters.

The 2021 Hilti Challenge dataset [9] used similar sensors to our proposed dataset and featured an AlphaSense 5-camera module which offered a 270° continuous field of view. For wide-coverage lidar measurements, [9] used the Ouster OS0-64 lidar with a range accuracy of ± 3 cm, whereas the dataset presented in this paper uses the much more accurate Hesai PandarXT-32 lidar with ± 1 cm range accuracy.

Compared to the aforementioned works, our dataset covers a variety of scenarios ranging from a sparse construction site to a 17th-century theatre with challenging staircases and narrow hallways. Throughout these indoor and outdoor sequences with difficult trajectories, we consistently provide millimeter-accurate ground truth positions at select control points using the method described in Sec. V-B. The challenging sequences and accurate sparse ground truth measurements of our dataset aimed to propel SLAM research into real-world applications such as architectural inspection and construction monitoring, where use cases require sub-centimeter accuracy.

III. HARDWARE

Our handheld multi-camera lidar inertial device, called Phasma, is shown in Fig. 1. Phasma is composed of these three sensors, rigidly assembled in a case that has been produced

Sensor	Type	Rate	Characteristics
Lidar	Hesai, PandarXT-32	10 Hz	32 Channels, 120 m Range 31° Vertical FoV
Cameras	Alphasense	40 Hz	1024 Horizontal Resolution 5 Global shutter (Infrared)
IMU	Bosch BMI085	400 Hz	720×540 pixels Synchronized with cameras

TABLE I: Overview of the sensors on the Phasma device.

by a precise milling machine. A complete URDF model of the device is also available as an open source ROS package³. The Hesai lidar is directly mounted below the cameras for a balanced and compact design so that the upward facing camera is not obstructed. The metal needle tip at the bottom is used to align with target points when determining ground truth (see Sec. V-B). Tab. I provides an overview of the sensor specifications.

Drawing on Hilti’s expertise, we manufactured a precise handheld device that had improved accuracy and stability compared the rig used in the previous challenge [9]. Specifically, by using milled components and dowel pins, we ensured an accurate assembly of the sensors. The correct placement of the lidar, metal tip, and cameras were verified using a GOM Atos Q3⁴ industrial 3D scanner. The multi-camera sensor is an Alphasense Core Development Kit from Sevensense Robotics AG. An FPGA within the Alphasense hardware synchronizes the IMU and five grayscale fisheye cameras – a frontal stereo pair with an 11 cm baseline, two lateral cameras, and one upward-facing camera. Each camera has a Field of View (FoV) of 126×92.4° and a resolution of 720×540 px. This configuration produces an overlapping FoV between the front and side cameras of about 36°. The cameras and the embedded cellphone-grade IMU operates at 40 Hz and 400 Hz, respectively. The Hesai lidar has 32 beams and a 31° elevation FoV, with a range of 5 cm to 120 m. Notably, the Hesai Pandar has a range accuracy of ± 1 cm and a precision of 0.5 cm (1σ).

A. IMU Calibration

A 90 min sequence of IMU data was collected on a stationary flat surface. We adopted the Allan Variance estimation method⁵ to compute the angle random walk, bias instability, and random walk for the gyroscope and the velocity random walk, bias instability, and random walk for the accelerometer. This IMU rosbag is provided with the dataset for the user’s convenience.

B. Camera Calibration

Similarly to the Newer College Dataset Multi-Camera Extension [10], we used the open source camera and IMU calibration toolbox Kalibr [11] to compute the intrinsic and extrinsic calibration of the Alphasense cameras. The calibration used the pinhole projection model with equidistant distortion. We then performed spatio-temporal calibration between the cameras and the IMU embedded in the Alphasense. The large

³<https://github.com/Hilti-Research/Hilti-SLAM-Challenge-2022>

⁴<https://www.gom.com/en/products/3d-scanning/atos-q>

⁵https://github.com/ori-drs/allan_variance_ros



Fig. 2: Hilti-Oxford dataset environments: *Top*: Construction site. *Middle*: Long office corridor (left), Sheldonian stairs (right). *Bottom*: Sheldonian lower gallery (left) and exterior (right).

rigid calibration target contained 7×12 April tags, with a tag size of 15 cm. All cameras were calibrated with the IMU, with the front cameras calibrated as stereo cameras and the remaining three calibrated as monocular cameras. In this dataset, we provide the rosbag of the camera and IMU calibration sequence to enable users to conduct their own calibration.

IV. DATASET

This dataset was recorded in two locations. The first is a construction site in Schaan, Liechtenstein, near the Hilti headquarters. It is a live construction site with limited texture and color variation. The site covers an area of $100\text{m} \times 30\text{m}$ with longer range measurements scanning nearby buildings. The site has four floor levels including a basement. The second location is the Sheldonian Theatre, built in 1664 in Oxford, England. The Sheldonian Theatre is used for ceremonial events and graduations and is an architecturally significant *listed building*. As shown in Fig. 2, the building spans 6 floors, from the basement to the octagonal cupola at the top. The cupola is accessible through narrow staircases, only 60 cm across. Both of these locations challenge SLAM systems in different ways.

Each dataset sequence is a rosbag that contains five camera image topics, one lidar topic, and one IMU topic. An example of the data is shown in Fig. 3. Below is a summary of each sequence. We qualitatively indicate the difficulty levels based on the environment and motions for each sequence as either easy, medium, or hard. Users can find more information including the top-down trajectories on the dataset website.

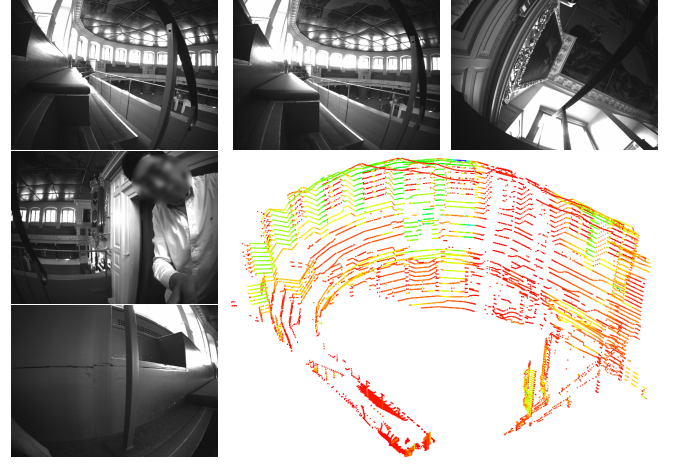


Fig. 3: Dataset example showing images from each camera, and a lidar scan of the upper gallery.

A. Challenge Sequences

Hilti Construction Site:

- a) **Exp01 Construction Ground Level** (Easy, 227 s): One loop around the ground level at typical walking speed.
- b) **Exp02 Construction Multilevel** (Medium, 430 s): One loop around the upper level then going down the staircase to the ground level. Some shaking motions with the angular velocity up to 3.5 rad/s
- c) **Exp03 Construction Stairs** (Hard, 309 s): Starting in the staircase, moving into dark corners while going down the stairs, then entering a car park in the basement, and finally returning to the top of the staircase. An operator is walking in front about half of the time.
- d) **Exp07 Long Corridor** (Medium, 132 s): A 100 m long corridor in Hilti's head office of 2 m width and 3 m height. Structurally, both ends of the corridor are higher than the middle section.

Oxford Sheldonian Theatre:

- a) **Exp09 Cupola** (Hard, 367 s): From the ground floor hall going up multiple levels through very narrow staircases to the very top of the theatre, the cupola. Then descending down another set of stairs back to the ground floor hall.
- b) **Exp11 Lower Gallery** (Medium, 151 s): From the ground floor hall to the first floor lower gallery, circling the lower gallery and back to the starting point.
- c) **Exp15 Attic to Upper Gallery** (Hard, 260 s): From the top floor attic space walking down to the upper gallery, going through some degenerate narrow spaces. Circling around the upper gallery and climbing back up another set of stairs to the attic space.
- d) **Exp21 Outside Building** (Easy, 152 s): Starting outside the theatre, circling the theatre and the main quad, and entering back into the theatre.

These eight sequences formed the 2022 Hilti Challenge.

B. Additional Sequences

We also release some additional sequences, with both sparse and dense ground truth trajectories, which can provide extra challenges for algorithm testing.

- a) **Exp04 Construction Upper Level 1** (Easy, 124 s):
One loop in the upper level with a typical walking speed.
- b) **Exp05 Construction Upper Level 2** (Easy, 125 s):
A repeat of Exp04 by another operator, offering a different walking pattern.
- c) **Exp06 Construction Upper Level 3** (Medium, 150 s):
Similar to 4 and 5 but a faster walking speed with aggressive motions, such as spinning on the spot, and approaching walls and corners.
- d) **Exp10 Cupola 2** (Hard, 446 s):
Similar to Exp09 but with faster motions.
- e) **Exp14 Basement 2** (Medium, 73 s):
A short sequence that starts from the staircase and goes through the basement with door opening scenarios.
- f) **Exp16 Attic to Upper Gallery 2** (Hard, 198 s):
Similar to Exp15 but with faster motions.
- g) **Exp18 Corridor Lower Gallery 2** (Hard, 100 s):
A short sequence starts in the corridor and enters the lower gallery from a different set of stairs from Exp11.
- h) **Exp23 The Sheldonian Slam** (Hard, 1049 s):
A large mission around the whole Sheldonian Theatre that visits all spaces and revisits the ground hall several times for loop closure purposes.

In particular, Exp10 and Exp16 are harder than the challenge sequences due to faster motions, and Exp23 is an interesting sequence which is much longer than the others and includes many loop closures, stair climbs, and sensor deprivations.

C. Characteristics of the Sequences

We intentionally introduce challenging and degenerate scenarios into the dataset. These scenarios include aggressive motions, such as shaking and swinging the device, dynamic objects occasionally blocking the field of view, narrow staircases which are geometrically similar, and dark corners where cameras cannot detect features.

The purpose of this is to test the robustness and accuracy of state-of-the-art SLAM systems. For example, Fig. 4 shows: (a) a person walking in front and blocking the field of view; (b) the handheld device being placed down close to one wall with only a few lidar points sensed in the environment (a degenerate mode for lidar-based SLAM); (c) the device entering a very dark corner in the basement, returning few lidar points. This is challenging for both vision- and lidar-based SLAM. In general, lidar-based SLAM suffers in confined space when there are not enough geometric constraints in a scan for registration. For vision-based SLAM, moving around in confined spaces can also result in rapid scene changes and fast image feature flow. We spent an average of 65 % of the time moving inside various confined spaces in Exp03 Construction Stair, Exp09 Cupola, and Exp15 Attic to Upper Gallery sequences.

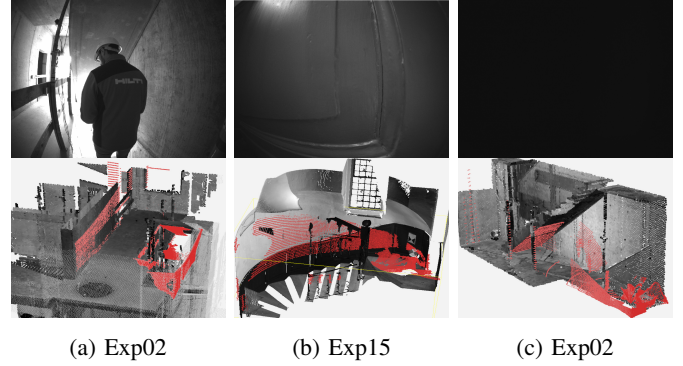


Fig. 4: Top row shows camera images in challenging scenarios with their corresponding lidar scans in red at the bottom (aligned to our ground truth model in grey).

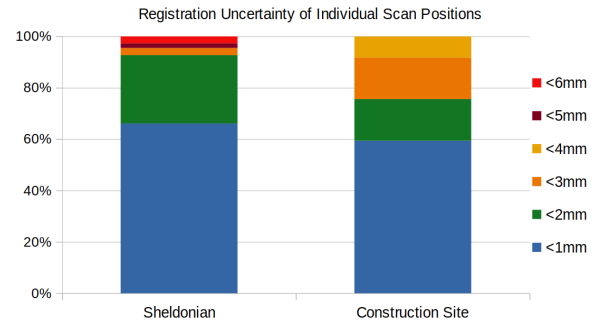


Fig. 5: Prior map individual scan registration uncertainties with respect to the initial scan.

V. GROUND TRUTH

A. Prior Map

Prior maps of the two facilities were built using the scanner shown in Fig. 6. The Z+F Imager 5016 3D laser scanner is equipped with an integrated HDR camera, internal light, and positioning system. It measures up to 360 m, with a maximum measurement rate of 1 million points/s. It has a field view of $360 \times 320^\circ$, an angular accuracy of 14.4 arcsec (both horizontal and vertical) and a linearity error of the laser system of $\leq 1 \text{ mm} + 10 \text{ ppm/m}$. The ranging noise is negligible (sub-mm). This scanner allows us to build accurate maps of the environment and establish ground truth points with millimeter accuracy. For the registration of the scans, we used reflective scanner targets as well as plane-to-plane registration followed by block adjustment [12]. Due to sufficient overlap between the laser scans, the reported pairwise registration uncertainty sigma is in the sub-mm range, which is the prior used before the bundle adjustment. The final uncertainty sigma for each scan with respect to the starting scan or master scan is shown in Fig. 5. 91 % and 95 % of scans have position uncertainty within 3mm for the Sheldonian and the construction site respectively. The most significant uncertainty is still just a few mm, as they are the leaf scan with fewer connections in the bundle adjustment graph. More importantly, we have not placed ground truth targets in those leaf scans. Some snapshots of the final registered point clouds are shown in Fig. 7.



Fig. 6: The Z+F scanner in the Sheldonian Theatre, scanning and taking images.

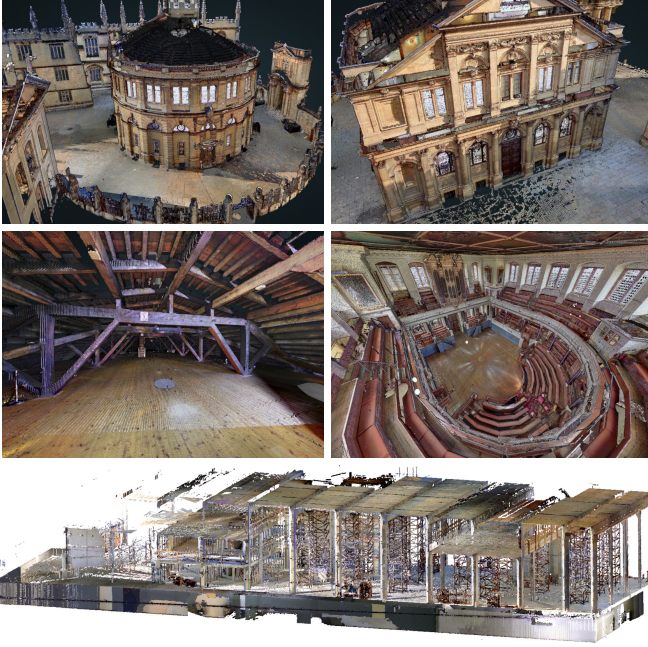


Fig. 7: Final point clouds built with the survey grade scanner. Top: Sheldonian theatre front and back exterior view. Mid: Sheldonian theatre attic and hall view. Bottom: Multi-level construction site side view.

B. Sparse Ground Truth For Evaluation

The process to set up a ground truth point is detailed as follows. We first select appropriate locations to set up crosshairs on the floor, by drawing on an adhesive blue marker 8. We then align the metal tip of a checkerboard target (shown in Fig. 8) at the center of the cross hairs. After adjusting the bubble level, each target is labelled numerically. These targets are scanned by the Z+F scanner and used in the fine registration step to create the complete point cloud. We ensure each target can be seen in multiple scans to add additional constraints to the cloud registration. All crosses drawn on the floor can be extracted from the final registered point cloud and used as ground truth evaluation points. During the handheld data gathering stage, we again place the tip of the handheld device at the crosshairs on the floor. To be noted, during this process, the site was closed to visitors and we ensured each blue marker stayed in the same place. Each time we took several seconds to carefully placed the metal tip on the crosshair, to ensure the manual

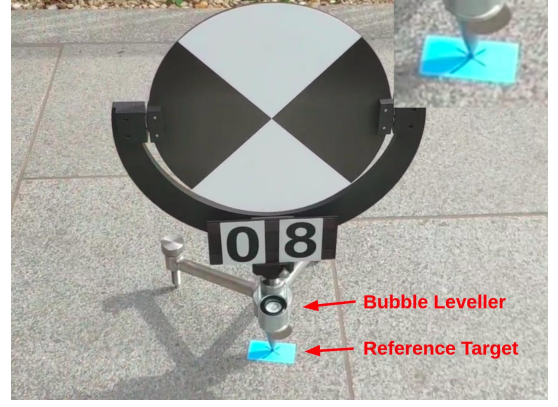


Fig. 8: A reference target placed on the ground and used for point cloud fine registration. Sparse ground truth points were created by placing Phasma at the cross-hairs of the reference target.

error stays less than 1 mm.

This sparse ground truth method is inspired by site surveying practice where the surveyor takes measurements of construction sites as a continuous monitoring or inspection procedure. We adopted the same equipment to establish ground truth poses for trajectories which is novel for SLAM datasets. It achieves millimeter accuracy but is limited to a small number (between 5 and 10) of instants where the Phasma device has to be laid on the crosshairs. To assist development and evaluation, denser but less accurate ground truth trajectories are provided for the additional sequences.

C. Dense Ground Truth For Development

The dense ground truth poses generation process uses the same method described in [10], [13]. The registered point cloud from the Z+F laser scanner in Sec. V-A was used as a prior map in which to localize. Ground truth poses are determined by registering each undistorted lidar scan to the prior map using an approach based on the point-to-point Iterative Closest Point method. We use our existing VILENS system [6] which is a Lidar-inertial odometry system to process each scan at a lower playback speed. Instead of building a map, we register to the prior ground truth map. We use a factor graph in the pre-integration IMU factor to motion correct the lidar scan. In general, the accuracy of this method is in the 1-2 cm range, so it can be useful to help users develop their SLAM algorithm, such as identifying precisely where odometry drift has occurred. When a SLAM algorithm approaches <1 cm accuracy, we recommend using the sparse ground truth in Sec. V-B for performance evaluation.

D. Scoring Metric

The aim of the challenge was to understand the state-of-the-art SLAM algorithms for use in the built environment, and in particular for the construction industry. Many real-world applications (e.g. autonomous hole drilling) require sub-centimeter accuracy to be useful. This motivated the accuracy-based error metric described below.

First, the control points and the estimated trajectory are aligned using SE(3) Umeyama alignment to account for any differences in coordinate systems [14]. Then the absolute distance error e_i between the i th control point and the estimated trajectory is calculated and given a score s_i , where

$$s_i = \begin{cases} 10 & \text{if } e_i < 0.01 \text{ m} \\ 6 & \text{if } e_i \in [0.01, 0.03 \text{ m}) \\ 3 & \text{if } e_i \in [0.03, 0.06 \text{ m}) \\ 1 & \text{if } e_i \in [0.06, 0.10 \text{ m}) \\ 0 & \text{if } e_i \geq 0.10 \text{ m} \end{cases} \quad (1)$$

The total score for each dataset S_j is the percentage of the maximum possible score (*i.e.* if all points had <1 cm error and scored 10 points),

$$S_j = \left(\frac{1}{10N} \sum_{i=0}^N s_i \right) \times 100 \quad (2)$$

where N is the number of ground truth points evaluated in each dataset. This denominator normalizes the score for a particular run to be between 0 and 100, regardless of the number of ground truth points in each dataset. The final score is then a sum of the scores from each experiment.

VI. CHALLENGE RESULTS AND FINDINGS

A. Results

A total of 42 academic and industrial groups submitted their results to the 2022 Hilti SLAM Challenge. The challenge results were announced as part of the Future of Construction workshop⁶ at the IEEE International Conference on Robotics and Automation in June 2022. To support their submissions, teams were given access to the calibration datasets, as well as three of the additional sequences with dense ground truth poses (see Sec. V-C).

The challenge results are shown in Tab. II. Summary reports of each team's approach are available on the challenge website⁷. The highest scoring team was CSIRO with a score of 563.8. Their Wildcat SLAM [15] algorithm uses a continuous-time trajectory representation for lidar-inertial odometry using sliding-window optimization and online pose graph optimization. This is refined by an offline global optimization module that takes advantage of non-causal information. Of the top 25 teams, all used lidar and IMU data, while only 10 used camera data.

The highest scoring vision-only submission was Smart Robotics Lab with a score of 32.5. SRL's OKVIS2.0 [16] produced complete trajectories for all of the sequences, however typical errors were in the 10–20 cm range which resulted in a low score. This highlights the gap in performance between lidar and camera-based SLAM systems, and the susceptibility of vision-based systems to subtle scaling and calibration errors.

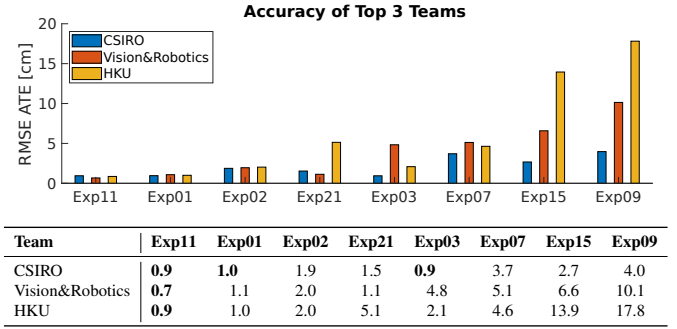


Fig. 9: Summary of the top three team's RMSE ATE (cm), sorted from smallest to largest error. Results in bold have reached the desired sub-cm accuracy.

B. Discussion

Tab. II describes the details of each algorithm's odometry and SLAM modules. In the odometry column, we classify each algorithm as either filter or optimization based and whether the odometry system is real-time. In the SLAM column, we label each algorithm according to the use of global bundle adjustment (Global BA), loop closure detection (LC), and if the algorithm only uses past measurements (causal). Overall, the scoring metric presented in this paper approximately aligns with the mean RMSE of the Absolute Trajectory Error (ATE). However, some entries, such as KTH&NTU, achieved a better ATE but lower score. This is explained by the fact that most of the poses fell outside the high-scoring sub-3 cm range. Additionally, the *mean* ATE can be deceptive here, as some teams have incomplete trajectories or performed badly in one particular sequence.

Fig. 9 shows a summary of the error on each sequence by the top three teams. Sequences with large open spaces and overlapping areas, where LIDAR scan matching can be highly effective, had the lowest error (Exp01, Exp02, Exp11, Exp21). The sequences with the highest error had challenging geometries for lidar-based algorithms including long narrow corridors (Exp07) and small staircases (Exp03, Exp09, Exp15), as illustrated in Fig. 2. This demonstrates that while the top three teams achieved accuracy close to 1 cm in some of the easier sequences, there is room for improvement in the others.

Another key observation is that the top four solutions were lidar-inertial only solutions, without the use of camera data. It has become common knowledge to fuse IMU measurements to provide a strong prior in SLAM system nowadays. While we still intended to create ill-conditioned situations for lidar-inertial based SLAM to require camera data fusion, the lidar scanner still captured sufficient information to avoid degeneracy in most sequences. For example, in Exp02 (Fig. 4(a)) we had an operator walk in front to block the sensors but the lidar scan was able to capture the slanted staircase ceiling and wooden rails (circled in Fig. 10-Right) which was enough to constrain the estimation. Similarly, when climbing the lower staircases of the Sheldonian, the lidar could scan through small windows onto adjacent buildings and the ground outside the theatre, despite there being insufficient constraints within the small staircase. These seemingly challenging situations were

⁶Video: https://www.youtube.com/watch?v=NpfJV_Q_SmK&t=205s

⁷<https://hilti-challenge.com/leader-board-2022.html>

	Lead Organization	Algorithm	Sensors Used			Odometry Type	Real-Time	SLAM			Same Params	Results	
			Lidar	IMU	Cam. (#)			Global BA	Causal	LC		ATE	Score
1	CSIRO	Wildcat SLAM [15]	✓	✓		SW Opt.	✓	✓	×	✓	✓	2.07	563.8
2	Vision & Robotics	MC2SLAM [17]	✓	✓		SW Opt.	✓	✓	×	✓	✓	3.94	443.8
3	HKU	FastLIO2[18], BALM [19]	✓	✓		Filter	✓	✓	–	✓	–	5.94	400.4
4	KAIST	Based on [18], [20]	✓	✓		Filter	–	✓	×	✓	–	19.02	317.5
5	Beihang Uni.	Based on [18], [21]	✓	✓	✓(2)	Filter	✓	×	✓	×	×	22.59	311.6
6	Luxembourg Uni.	Based on [18], [22]	✓	✓	✓(1)	Filter	✓	×	✓	×	✓	20.49	303.8
7	MINES ParisTech	CT-ICP [23]	✓	✓		Opt.	✓	×	✓	✓	–	7.72*	272.8
8	AIST	VITAMIN-E [24], [25]	✓	✓	✓(3)	SW Opt.	✓	×	✓	×	✓	16.16	260.5
9	HKUST & Georgia Tech	Based on [18], [26]	✓	✓	✓(5)	Filter	✓	✓	×	✓	✓	47.50	257.6
10	KTH & NTU	VIRAL SLAM [27]	✓	✓		SW Opt.	×	×	✓	×	✓	6.90	251.9
Vision-only Results													
1	TUM	OKVIS2.0 [16]		✓	✓(5)	SW Opt.	✓	✓	×	✓	✓	25.36	32.5
2	Stuttgart Uni. & TUM	Based on [28]		✓	✓(4)	SW Opt.	×	×	✓	×	×	42.04	22.2

Legend: # = No. of cameras used, – = No information provided, ATE = Mean RMSE ATE (cm), * = Did not submit results for Exp15

TABLE II: HILTI SLAM Challenge 2022 Results

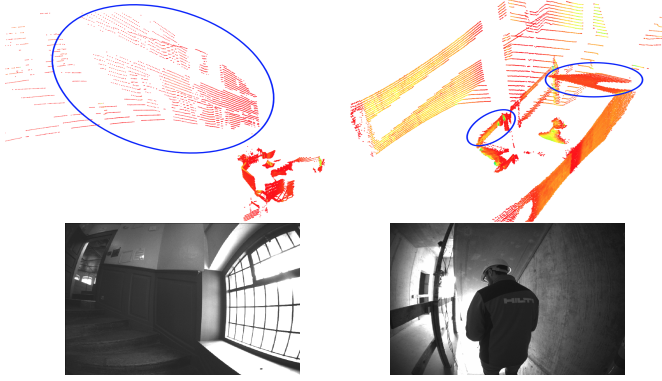


Fig. 10: *Left:* While in a narrow staircase, the lidar sees through a window and scans the adjacent building and the ground. *Right:* An operator walks in front to block the view, but the stair ceiling and rails are scanned and provide enough constraints.

resolved by the accuracy and range of the lidar. We note that the top performing teams relied on dense local lidar submaps to overcome these locally-degenerate scenarios.

The most challenging sequence was Exp09 which entered the narrow upper staircases of the Sheldonian (from 132 s) where there were no windows, providing limited constraints for lidar odometry. Other failure situation example are dark corners under the stair cases (Exp03, 71 s), narrow space (Exp15, 24 s), middle of a long corridor (Exp07, 46 s).

There are some interesting findings on loop closure from the submissions. First, for shorter sequences, e.g. Exp01, teams could simply keep a complete map in their lidar odometry system. Drift was small enough that an explicit loop closure was not actually needed, and the system could implicitly localise to the local map without doing pose graph SLAM. For longer runs, such as Exp03, loop closures were used to but sometimes they distributed the error from one particular section to the whole trajectory. More importantly, teams used loop closures between sequences to create a multi-sequence map, e.g. linking exp09, exp11 and exp15, to further reduce drift.

C. Known Issues

Although we took the utmost care in the creation of this dataset and challenge, there were a few limitations of our

approach:

Scoring Metric: The metric used in this evaluation focused on the accuracy of the trajectory and did not consider other performance characteristics of real-time SLAM systems, such as latency and computation. There was a wide range of timing and computational differences for various online and offline processing methods. Some teams used multi-sequence fusion in post-processing to optimize their results - thus their results are likely to be better than when operating on the individual sequences. Also given each team used different hardware, it would be difficult to include this in the scoring. However, we have qualitatively captured these traits in Table II and would refer readers to each team's report on the website.

Lidar-IMU Calibration: While the dataset used highly-accurate, machined components with low tolerances, we did not undertake a separate extrinsic calibration between the lidar and IMU. This may have resulted in a few millimeters of error in the ground truth control points.

VII. CONCLUSION

This paper presents the Hilti-Oxford Dataset, comprised of vision, lidar, and inertial sensing collected in two challenging environments. The provision of highly accurate ground truth enables the transparent evaluation of SLAM systems. In this challenge, we found that the top three teams achieved 3 cm accuracy in the construction site sequences but incurred higher errors for the harder sequences from the Sheldonian.

The HILTI SLAM Challenge leaderboard remains live and can accept new submissions for automatic evaluation. We hope this dataset provides researchers with a difficult and diverse challenge to improve their SLAM systems.

VIII. ACKNOWLEDGEMENTS

We thank the organizers of the Future of Construction workshop at ICRA 2022 for hosting the challenge and Beda Berner from HILTI who produced and calibrated the handheld device used to collect this dataset. M. Fallon thanks the Royal Society for funding his University Research Fellowship.

REFERENCES

- [1] D. Schubert, T. Goll, N. Demmel, V. Usenko, J. Stückler, and D. Cremers, "The TUM VI Benchmark for Evaluating Visual-Inertial Odometry," in *IEEE/RSJ Intl Conf. on Intelligent Robots and Systems*, 2018, pp. 1680–1687.
- [2] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The EuRoC micro aerial vehicle datasets," *Int. Journal of Robotics Res.*, vol. 35, no. 10, pp. 1157–1163, 2016.
- [3] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. Journal of Robotics Res.*, vol. 32, no. 11, 2013.
- [4] C. Debeunne and D. Vivet, "A Review of Visual-LiDAR Fusion based Simultaneous Localization and Mapping," *Sensors*, vol. 20, no. 7, 2020.
- [5] 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, pp. 1182–1189, 2022.
- [6] D. Wisth, M. Camurri, and M. Fallon, "VILENS: Visual, Inertial, Lidar, and Leg Odometry for All-Terrain Legged Robots," *IEEE Trans. on Robotics*, 2022.
- [7] S. Yogamani, C. Hughes, J. Horgan, G. Sistu, S. Chennupati, M. Uricar, S. Milz, M. Simon, K. Amende, C. Witt, H. Rashed, S. Nayak, S. Mansoor, P. Varley, X. Perrotton, D. Odea, and P. Pérez, "WoodScape: A Multi-Task, Multi-Camera Fisheye Dataset for Autonomous Driving," in *IEEE/CVF Intl. Conf. on Computer Vision*, 2019, pp. 9307–9317.
- [8] N. Carlevaris-Bianco, A. K. Ushani, and R. M. Eustice, "University of Michigan North Campus long-term vision and lidar dataset," *Intl. Journal of Robotics Research*, vol. 35, no. 9, pp. 1023–1035, 2016.
- [9] M. Helmberger, K. Morin, B. Berner, N. Kumar, G. Cioffi, and D. Scaramuzza, "The Hilti SLAM Challenge Dataset," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7518–7525, 2022.
- [10] L. Zhang, M. Camurri, D. Wisth, and M. Fallon, "Multi-Camera LiDAR Inertial Extension to the Newer College Dataset," *arXiv:2112.08854*, 2021.
- [11] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending Kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes," in *IEEE Intl. Conf. on Robotics and Automation*, 2016, pp. 4304–4311.
- [12] D. Wujanz, S. Schaller, F. Gielsdorf, and L. Gruendig, "Plane-based registration of several thousand laser scans on standard hardware," *Photogrammetry, Remote Sensing and Spatial Info. Sciences*, vol. XLII-2, pp. 1207–1212, 2018.
- [13] M. Ramezani, Y. Wang, M. Camurri, D. Wisth, M. Mattamala, and M. Fallon, "The Newer College Dataset: Handheld LiDAR, Inertial and Vision with Ground Truth," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2020, pp. 4353–4360.
- [14] M. Grupp, "evo: Python package for the evaluation of odometry and SLAM." <https://github.com/MichaelGrupp/evo>, 2017.
- [15] M. Ramezani, K. Khosoussi, G. Catt, P. Moghadam, J. Williams, P. Borges, F. Pauling, and N. Kottege, "Wildcat: Online Continuous-Time 3D Lidar-Inertial SLAM," *arXiv:2205.12595*, 2022.
- [16] S. Leutenegger, "OKVIS2: Realtime Scalable Visual-Inertial SLAM with Loop Closure," *arXiv preprint arXiv:2202.09199*, 2022.
- [17] F. Neuhaus, T. Koß, R. Kohnen, and D. Paulus, "MC2SLAM: Real-Time Inertial Lidar Odometry Using Two-Scan Motion Compensation," *German Conference on Pattern Recognition*, pp. 60–72, 2019.
- [18] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "FAST-LIO2: Fast Direct LiDAR-Inertial Odometry," *IEEE Trans. on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [19] Z. Liu and F. Zhang, "BALM: Bundle Adjustment for Lidar Mapping," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3184–3191, 2021.
- [20] F. A. Maken, F. Ramos, and L. Ott, "Estimating Motion Uncertainty with Bayesian ICP," in *IEEE Intl. Conf. on Robotics and Auto.*, 2020.
- [21] T. Qin, P. Li, and S. Shen, "VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator," *IEEE Trans. on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [22] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, "OpenVINS: A Research Platform for Visual-Inertial Estimation," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2020, pp. 4666–4672.
- [23] P. Dellenbach, J.-E. Deschaud, B. Jacquet, and F. Goulette, "CT-ICP: Real-time Elastic LiDAR Odometry with Loop Closure," in *IEEE Intl Conf on Robotics and Automation*, 2022, pp. 5580–5586.
- [24] M. Yokozuka, S. Oishi, S. Thompson, and A. Banno, "VITAMIN-E: Visual Tracking and Mapping With Extremely Dense Feature Points," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2019.
- [25] K. Koide, M. Yokozuka, S. Oishi, and A. Banno, "Globally Consistent and Tightly Coupled 3D LiDAR Inertial Mapping," in *IEEE Intl. Conf. on Robotics and Automation*, 2022, pp. 5622–5628.
- [26] F. A. Maken, F. Ramos, and L. Ott, "Estimating motion uncertainty with Bayesian ICP," in *IEEE Intl. Conf. on Robotics and Automation*, 2020, pp. 8602–8608.
- [27] T. Nguyen, S. Yuan, M. Cao, T. H. Nguyen, and L. Xie, "VIRAL SLAM: Tightly Coupled Camera-IMU-UWB-Lidar SLAM," *CoRR*, vol. abs/2105.03296, 2021.
- [28] Z. Teed and J. Deng, "DROID-SLAM: Deep Visual SLAM for Monocular, Stereo, and RGB-D Cameras," *Advances in Neural Information Processing Systems*, vol. 34, pp. 16 558–16 569, 2021.