

The Oxford Spires Dataset: Benchmarking Large-Scale LiDAR-Visual Localisation, Reconstruction and Radiance Field Methods

Journal Title
XX(X):1–17
©The Author(s) 2025
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Yifu Tao^{1*}, Miguel Ángel Muñoz-Bañón^{1,2*}, Lintong Zhang¹, Jiahao Wang¹, Lanke Frank Tarimo Fu¹, and Maurice Fallon¹

Abstract

This paper introduces a large-scale multi-modal dataset captured in and around well-known landmarks in Oxford using a custom-built multi-sensor perception unit as well as a millimetre-accurate map from a Terrestrial LiDAR Scanner (TLS). The perception unit includes three synchronised global shutter colour cameras, an automotive 3D LiDAR scanner, and an inertial sensor — all precisely calibrated. We also establish benchmarks for tasks involving localisation, reconstruction, and novel-view synthesis, which enable the evaluation of Simultaneous Localisation and Mapping (SLAM) methods, Structure-from-Motion (SfM) and Multi-view Stereo (MVS) methods as well as radiance field methods such as Neural Radiance Fields (NeRF) and 3D Gaussian Splatting. To evaluate 3D reconstruction the TLS 3D models are used as ground truth. Localisation ground truth is computed by registering the mobile LiDAR scans to the TLS 3D models. Radiance field methods are evaluated not only with poses sampled from the input trajectory, but also from viewpoints that are from trajectories which are distant from the training poses. Our evaluation demonstrates a key limitation of state-of-the-art radiance field methods: we show that they tend to overfit to the training poses/images and do not generalise well to out-of-sequence poses. They also underperform in 3D reconstruction compared to MVS systems using the same visual inputs. Our dataset and benchmarks are intended to facilitate better integration of radiance field methods and SLAM systems. The raw and processed data, along with software for parsing and evaluation, can be accessed at <https://dynamic.robots.ox.ac.uk/datasets/oxford-spires/>.

Keywords

Dataset, Localisation, 3D Reconstruction, Novel-View Synthesis, SLAM, NeRF, Radiance Field, LiDAR Camera Sensor Fusion, Colour Reconstruction, Calibration

1 Introduction

Localisation and 3D reconstruction are fundamental problems in both robotics and computer vision, with applications spanning autonomous driving, building inspection and augmented reality. There are methods that focus on localisation (e.g. visual or lidar odometry, Structure-from-Motion (SfM), place recognition and relocalisation) and others that focus on 3D reconstruction/mapping (e.g. Multi-view Stereo (MVS) and occupancy mapping). In mobile robotics, both problems can be solved concurrently by Simultaneous Localisation and Mapping (SLAM) methods, which are our primary focus in this work. For large-scale outdoor environments, cameras and LiDARs are the most commonly used exteroceptive sensor modalities for environment perception tasks. Additionally, the Inertial Measurement Unit (IMU) is a common interoceptive sensor. The two mentioned exteroceptive sensor technologies have complementary characteristics: LiDAR captures long-range depth measurements that are accurate but sparse, while camera images capture texture with higher resolution.

The evaluation of the outdoor SLAM systems has primarily focused on localisation accuracy, conversely

quantitative evaluation of the 3D reconstruction quality is often lacking. An important reason for this is the limited availability of high-quality ground truth. For 3D reconstruction, ground truth is typically collected using survey-grade Terrestrial LiDAR Scanners (TLS) which are expensive (Zhang et al. 2022). Compared to indoor scenes where MVS and RGB-D SLAM systems are often evaluated, the large-scale of outdoor environments makes TLS data collection laborious. As a result, many outdoor SLAM datasets do not include precise ground truth reconstruction from TLS and rely on ground truth trajectories from other sensors such as GNSS-RTK (Geiger et al. 2013).

In addition to geometric reconstruction, colour reconstruction is becoming more important with the advances of

¹Oxford Robotics Inst., Dept. of Eng. Science, Univ. of Oxford, UK

²Group of Automation, Robotics and Computer Vision (AUROVA), University of Alicante, Spain

*denotes equal contribution

Corresponding author:

Yifu Tao, Oxford Robotics Inst., Dept. of Eng. Science, Univ. of Oxford
OX1 3PJ, UK.

Email: oxfordspiresdataset@robots.ox.ac.uk




Figure 1. *Top:* Point cloud of the Radcliffe Camera and the Bodleian Library captured by the TLS. Note the scale bar in the bottom-right corner. *Bottom:* Views of the TLS maps from other sites. Each column is a different site. The upper images show views of colour point clouds while the lower images were taken by the scanner's cameras.

radiance field methods including Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) and 3D Gaussian Splatting (Kerbl et al. 2023). These methods take as input calibrated camera images and their precise 3D poses (typically estimated using SfM), and output a dense 3D field with volume density (similar to differential opacity) and view-dependent colour. The output radiance field can be used to synthesise photorealistic images using volume rendering techniques. Since radiance field methods are capable of representing complex geometry and appearance, some SLAM systems have adopted it as their underlying 3D map representation (Sucar et al. 2021; Zhu et al. 2022).

Despite the rapid development of radiance field methods, their use in outdoor mobile robot perception has been less well explored. Radiance field methods are often evaluated by the quality of images rendered using datasets where the image set points at a single object observed in controlled lighting conditions, and often indoors. For a mobile robot operating in an outdoor environment, the trajectory is usually not object-centric, and viewpoints are relatively sparse

compared to the size of the scene. Inferring 3D structure from monocular images alone is more challenging if provided with fewer viewpoint constraints, and this can lead to artefacts (e.g. the elongated Gaussians along the viewing direction mentioned in Matsuki et al. (2024a)) that are not noticeable if only evaluated from nearby poses. In addition, radiance field methods can also generate “floater” artefacts to overfit per-frame lighting conditions (as discussed in Tancik et al. (2023)) and the texture of the sky. Both are common challenges in outdoor environments. Such artefacts can lead to inferior 3D reconstruction and poorer photorealistic rendering from a pose that is far from the training sequence. To develop radiance field methods that can be integrated with outdoor SLAM systems, it is crucial to have a dataset with colour images, LiDAR and accurate ground truth trajectory and reconstruction.

In this work, we introduce the Oxford Spires Dataset, a large-scale dataset collected across six historical landmarks in Oxford, UK, covering more than 20 000 m² per site. The total area recorded in this dataset is more than 125 000 m²,

about the size of a small town. It provides high-resolution RGB image streams from three cameras, 3D wide Field-of-View LiDAR data, and inertial data from a mobile handheld device. It is accompanied by millimetre-accurate reference scans which serve as the reference ground truth 3D model for reconstruction systems. We also use it to determine the ground truth trajectories of the handheld device. The setup of three colour cameras facing forward, left and to the right is a particularly novel characteristic as similar datasets contain TLS-based ground truth, LiDAR scans but only forward-facing colour camera(s) (Wei et al. 2024; Nguyen et al. 2024; Liu et al. 2024). The side-facing cameras provide increased Field-of-View that not only increases the texture mapping coverage but also provides view constraints which are crucial for vision-only systems to infer the 3D structure. The three-camera setup also makes textured mapping more tractable from a simple linear pass through an environment, rather than requiring exhaustive scanning. This makes our dataset suitable for evaluating radiance field methods in outdoor mobile robotics and reality capture contexts. Leveraging this rich combination of sensor data, we also introduce three benchmarks for localisation, 3D reconstruction, and novel-view synthesis. We use the benchmark to evaluate state-of-the-art SLAM systems, SfM-MVS systems and radiance field systems. In particular, the novel-view synthesis benchmark features test data not only sampled from a single reference trajectory, but also from other sequences where the device travelled in an opposing direction and along a trajectory far from the reference. The evaluation results of state-of-the-art radiance field methods highlight the problem of overfitting to the training data and an inability to generalise to distant viewpoints. Our dataset opens up new research avenues in this space. We release the raw sensor data as well as processed data including example outputs from a LiDAR SLAM system (such as motion-undistorted point clouds) and a SfM system (which can be used for MVS, NeRF and 3D Gaussian Splatting), as well as ground truth trajectories and reconstruction. Software for parsing the data and evaluating the systems presented in the three benchmarks is also made available.

In summary, our main contributions are as follows:

- A large-scale outdoor dataset collected at six historical sites, covering an average area of about 10 000 m² each. In total, 24 sequences were recorded, with the average distance travelled in each sequence exceeding 400 metres.
- The dataset is collected with a sensor suite comprising three 1.6 megapixel global shutter fisheye RGB cameras, a wide Field-of-View 64-beam 3D LiDAR, and inertial data, paired with millimetre-accurate reference 3D models captured using a TLS.
- We provide precise calibrations for the synchronised sensors — including 3 fisheye cameras, IMU and the 104° FoV LiDAR.
- Three benchmarks for localisation, reconstruction and novel-view synthesis with ground truth generated using the 3D models from the TLS. In this paper, we evaluated state-of-the-art SLAM, SfM, MVS, NeRF and 3D Gaussian Splatting methods for each.
- Evaluation software is released for using the dataset and benchmarking methods.

2 Related Work

In this section, we overview related datasets that are available for evaluating localisation, 3D reconstruction and novel-view synthesis. A summary of these datasets is presented in Tab. 1.

2.1 Datasets for Evaluating Localisation

Localisation is a key task in robotics and computer vision, and is performed by methods including odometry, SLAM, SfM, place recognition or relocalisation in a prior map. In indoor environments, cameras and RGB-D sensors are commonly used. TUM RGB-D (Sturm et al. 2012) is one of the first benchmarks which sought to evaluate localisation performance using ground truth trajectories. While for visual-inertial SLAM systems, EuROC (Burri et al. 2016) and TUM VI (Schubert et al. 2018) are popular datasets used in the research community. For outdoor environments, LiDAR is a common sensor modality and has been used in robotics datasets such as New College (Smith et al. 2009) and NCLT (Carlevaris-Bianco et al. 2016). Other datasets focus on evaluating odometry and SLAM trajectories in the context of autonomous driving, including KITTI (Geiger et al. 2013), Complex Urban (Jeong et al. 2019), and WoodScape (Yogamani et al. 2019). Several datasets including RELIS-3D (Jiang et al. 2021), Botanic Garden (Liu et al. 2024), and WHU-Helmet (Li et al. 2023) focus on unstructured natural environments (forest and rural areas). Robot platforms were used to collect datasets including quadrupeds (Brizi et al. 2024; Wei et al. 2024) and aerial robots (Li et al. 2024).

To evaluate the accuracy of localisation systems, a precise ground truth trajectory is essential. For self-driving datasets, ground truth trajectories are often obtained by fusing GNSS data with inertial and LiDAR data (Geiger et al. 2013). One limitation of GNSS-based ground truth is that it is not reliable in areas such as urban canyons. Motion capture systems can also be used to obtain ground truth trajectories (Helmberger et al. 2022), although they are often limited to indoor environments. In outdoor environments, Newer College (Ramezani et al. 2020b) generate a centimetre-accurate ground truth by registering LiDAR scans against an accurate prior map obtained using TLS. The approach of registering mobile lidar scans to an accurate prior map was also adopted by the authors of BotanicGarden (Liu et al. 2024) and MCD (Nguyen et al. 2024). Hilti-Oxford (Zhang et al. 2022) is notable in achieving millimetre accuracy ground truth for a sample set of stationary poses by using reference targets. Our work follows the approach used in Newer College (Ramezani et al. 2020b) to generate dense ground truth trajectories.

2.2 Datasets for Evaluating 3D Reconstruction

SLAM systems estimate both a robot/sensor trajectory and the map of their environment; however, many SLAM datasets only provide ground truth trajectories. Few datasets evaluate the accuracy of the map reconstruction, because accurate ground truth reconstructions are costly and laborious to obtain. Because of this, some SLAM datasets such as ICL-NUIM (Handa et al. 2014) create ground truth 3D models using simulation, and other datasets including Matterport3D (Chang et al. 2017) and ScanNet (Dai et al.

Table 1. Summary of related datasets for testing localisation and reconstruction methods. Oxford Spires is a large-scale outdoor SLAM dataset with three colour camera images, LiDAR as well as ground truth trajectories and 3D models. It can be used to evaluate tasks including localisation, reconstruction and novel-view synthesis. Features in other related datasets that are not suitable to our target domain (outdoor SLAM with colour reconstruction) are coloured with a **grey background**. These features include indoor scenes, greyscale camera images, short range (<10 m) depth sensing, and imprecise or missing ground truth 3D models.

Dataset	Scene	Camera			Depth Sensor	Ground Truth	
		Colour	Shutter	Resolution		Trajectory	3D Model
New College (Smith et al. 2009)	Outdoor	RGB/Grey	-	0.2 MP	2 × LM2 291-S14	-	-
TUM-RGBD (Sturm et al. 2012)	Indoor	RGB	Rolling	0.3 MP	MS Kinect RGB-D	MoCap	-
KITTI (Geiger et al. 2013)	Outdoor	RGB/Grey	Global	1.4 MP	Velodyne HDL-64E	GPS RTK	-
NCLT (Carlevaris-Bianco et al. 2016)	Outdoor&Indoor	RGB	Global	1.9 MP	Velodyne HDL-32E	GPS RTK	-
EuROC (Burri et al. 2016)	Indoor	Greyscale	Global	0.4 MP	-	Leica&Vicon	Leica MS50
DTU (Aanæs et al. 2016)	Indoor	RGB	-	1.9 MP	-	Robot Arm	Structured light scanner
ScanNet (Dai et al. 2017)	Indoor	RGB	Rolling	1.3 MP	Structure Sensor RGB-D	RGB-D SLAM	RGB-D SLAM
ETH3D (Schops et al. 2017)	Outdoor&Indoor	RGB/Grey	Global	0.4/24 MP	-	COLMAP+ICP	FARO Focus X330
Tanks and Temples (Knapitsch et al. 2017)	Outdoor&Indoor	RGB	Rolling	8 MP	-	Mutual Information	FARO Focus X330
Complex Urban (Jeong et al. 2019)	Outdoor	RGB	Global	0.7 MP	Velodyne VLP-16	GPS+LiDAR SLAM	-
WoodScape (Yogamani et al. 2019)	Outdoor	RGB	Rolling	1 MP	Velodyne HDL-64E	GNSS-IMU	-
Newer College (Ramezani et al. 2020b)	Outdoor	Greyscale	Global	0.4 MP	Ouster OS1-64	ICP	Leica BLK360
Hilti-21 (Helmberger et al. 2022)	Outdoor&Indoor	Greyscale	Global	1.3MP	Ouster OS0-64	MoCap/Total station	-
Hilti-22 (Zhang et al. 2022)	Outdoor&Indoor	Greyscale	Global	0.4 MP	Hesai XT32	ICP+Reference Target	Z+F Imager 5016
LaMAR (Sarlin et al. 2022)	Outdoor&Indoor	RGB/Grey	Both	2.0 MP	HoloLens2/iPhone RGB-D	LiDAR SLAM+SfM	LiDAR SLAM+SfM
WHU-Helmet (Li et al. 2023)	Outdoor&Indoor	RGB	-	-	Livox Avia	LiDAR SLAM	LiDAR SLAM
RELLIS-3D (Jiang et al. 2021)	Outdoor	RGB	Global	2.3 MP	VLP-16&OS1-64	LiDAR SLAM	LiDAR SLAM
ROVER (Schmidt et al. 2024)	Outdoor	RGB/Gray	Both	0.7 MP	-	Visual SLAM	-
Hilti-23 (Nair et al. 2024)	Indoor	Greyscale	Global	1.0 MP	Robosense B Pearl	Reference Target	Trimble X7
ScanNet++ (Yeshwanth et al. 2023)	Indoor	RGB	Rolling	2.8 MP	iPhone RGB-D	COLMAP	FARO Focus Premium
Hilti-24 (Sun et al. 2023)	Indoor	-	-	-	Matterport RGB-D	ICP	Matterport RGB-D
MARS-LVIG (Li et al. 2024)	Outdoor	RGB	Global	5.0 MP	Livox Avia	GNSS RTK	LiDAR-GNSS Mapping
VBR (Brizi et al. 2024)	Outdoor&Indoor	RGB	Global	0.9 MP	Ouster OS0-128	LiDAR-GNSS	-
BotanicGarden (Liu et al. 2024)	Outdoor	RGB/Grey	Global	2.3 MP	VLP16+Livox Avia	ICP	Leica RTC360
FusionPortableV2 (Wei et al. 2024)	Outdoor&Indoor	RGB	Global	0.7 MP	Ouster OS1-128	Prism-GNSS-IMU	Leica RTC/BLK360
MCD (Nguyen et al. 2024)	Outdoor&Indoor	RGB	Global	1.0 MP	Ouster OS1 +Livox Mid70	ICP	Survey-grade map
Oxford Spires (ours)	Outdoor&Indoor	RGB	Global	1.6 MP	Hesai QT64	ICP	Leica RTC360

2017) actually use the output from a RGB-D SLAM system as ground truth. Replica (Straub et al. 2019) provides higher-quality 3D meshes than ScanNet and Matterport 3D, and the rendered images are photo-realistic. The RGB-D SLAM ground truth approach cannot be adapted to outdoor scenes due to the short range of depth cameras. LaMAR (Sarlin et al. 2022) includes outdoor sequences and a ground truth 3D model from a combination of VIO, SLAM and SfM. The ground truth 3D model obtained with SLAM is generally not as accurate as what TLS can produce. Survey-grade TLS achieves millimetre-level accuracy (A comparison can be found in ScanNet++ (Yeshwanth et al. 2023)).

Among the datasets that provide precise ground truth reconstruction (obtained from TLS), EuROC (Burri et al. 2016) and ScanNet++ (Yeshwanth et al. 2023) are captured from indoor environments, and hence LiDAR is not used. The only available outdoor SLAM datasets that include accurate ground truth 3D models (to the best of our knowledge) are Newer College (Ramezani et al. 2020b) and Hilti-Oxford-2022 (Zhang et al. 2022). Both datasets use relatively low-resolution greyscale cameras, and therefore are not suitable for colour 3D reconstruction. Compared to them, our dataset provides high-resolution colour images from three cameras, and is hence suitable for not only 3D reconstruction but also novel-view synthesis.

In the field of computer vision, datasets with accurate ground truth reconstruction exist for MVS research, but often they target small-scale indoor scenes. Middlebury (Seitz et al. 2006) was one of the early datasets with ground truth depth obtained using a structured light scanner. DTU (Aanæs et al. 2016) captured individual objects using a robotic arm in a controlled environment, with its ground truth

also obtained using structured light scans. ETH3D (Schops et al. 2017) provides both high-resolution images (<80 per sequence) recorded by a DSLR camera, and low-resolution synchronised grey-scale images (~1000 per sequence), with a ground truth 3D model obtained using a TLS. Tanks and Temples (Knapitsch et al. 2017) is another popular benchmark for 3D reconstruction with ground truth from TLS. It can be used to evaluate both SfM and MVS algorithms and uses a higher-quality camera for its video data.

2.3 Datasets for Evaluating Novel-view Synthesis

Radiance fields have emerged to the most promising representation for novel-view synthesis. The input images for radiance field methods are often co-registered using SfM methods such as COLMAP (Schönberger and Frahm 2016). In the original NeRF paper (Mildenhall et al. 2021), both synthetic datasets and real-world image sequences from LLFF (Mildenhall et al. 2019) were used. Subsequently, Mip-NeRF 360 (Barron et al. 2022) included object-centric framed images taken in indoor and outdoor environments, and is popular in the radiance field community. Radiance field methods are typically evaluated using test set images that are sampled from an input trajectory and excluded from training. ScanNet++ (Yeshwanth et al. 2023) used a more challenging evaluation approach. The authors capture test images independently from the training sequence using a higher-quality DSLR camera in the indoor environment. In contrast, our dataset focuses on large-scale outdoor environments, and provides test images captured from

sequences with distant viewpoints. In this manner, we aim to advance the generalisation capability of existing radiance field methods.

3 Hardware

3.1 Handheld Perception Unit

Our perception unit, called Frontier, has three cameras, an IMU and a LiDAR. It is shown in Fig. 2 and Fig. 3 (right). The three colour fisheye cameras face forward, left, and right from a customised Alphasense Core Development Kit from Sevensense Robotics AG. Each camera has a Field-of-View of $126^\circ \times 92.4^\circ$ with a resolution of 1440×1080 pixels. There is a cellphone-grade IMU in the Alphasense Core Development Kit, which is hardware-synchronised with the three cameras using an FPGA from Sevensense. The three cameras have around 36° overlap, which enables multi-camera calibration mentioned in Sec. 4.1. The cameras operate at 20 Hz, and the IMU operates at 400 Hz. Auto-exposure is enabled for the cameras to capture indoor and outdoor scenes with different lighting conditions. A 64-channel Hesai QT64 LiDAR operating at 10 Hz was mounted on top of the cameras, with a Field-of-View of 104° , and a maximum range of 60 m. The LiDAR’s accuracy is ± 3 cm (typical), and the relation between reflectance and intensity is $reflectance = \alpha \times intensity \times range^2$.

Multi-sensor synchronisation is achieved both on the hardware level and software level. In terms of software synchronisation, we synchronise the device clock of Alphasense Core Development Kit (with the three cameras and IMU) and the Hesai LiDAR to the host computer on the Frontier using the Precision Time Protocol (PTP) which achieves sub-microsecond accuracy. As for hardware synchronisation, the three cameras are synchronised with the IMU within the Alphasense Core Development Kit, and we synchronise a LiDAR point cloud with the cameras using motion undistortion. While the three cameras’ exposure times are not identical for one shutter cycle due to the auto-exposure, the exposure intervals of the cameras are aligned in the middle[†]. This means that the timestamps of the mid-frame of the exposure interval of each camera are identical. This timestamp is used as the timestamp for the camera image. The Hesai QT64 LiDAR is a rolling shutter sensor, and the point cloud data is not triggered at a specific timestamp but is obtained continuously. To obtain a synchronised LiDAR point cloud for a synchronised set of three camera images, we motion-correct a LiDAR point cloud with IMU preintegration using VILENS (Wisth et al. 2023). The point cloud is undistorted to the same time as the next camera image frame. More information regarding the motion undistortion can be found in Wisth et al. (2021). In summary, each node in the SLAM pose graph has three camera images and an undistorted LiDAR point cloud with identical timestamps, which can be used for 3D reconstruction and novel-view synthesis.

3.2 Millimetre-accurate TLS

To obtain an accurate 3D reference model to benchmark localisation and reconstruction, we used a Leica RTC360 TLS (Fig. 3, left). It has a maximum range of 130 m and a




Figure 2. An isometric view of the sensor setup highlighting the coordinate frames of the cameras, the IMU and the LiDAR.



Figure 3. Leica RTC360 TLS and the Frontier device in Blenheim Palace (left) and Christ Church College (right).

Field-of-View of $360^\circ \times 300^\circ$. The final 3D point accuracy is 1.9 mm at 10 m and 5.3 mm at 40 m. The point clouds are coloured using 432 mega-pixel images captured by three cameras. Scans are registered in the field and re-optimised later using Leica’s Cyclone REGISTER 360 Plus software. The average cloud-to-cloud error in our sites ranges from 3-7 mm. After merging all the scans, the resultant point cloud is very large (10GB), so for ease of use we downsampled the TLS scan to 1cm resolution. Nonetheless, we provide the original raw TLS scans in our dataset.

4 Calibration

4.1 Multi-Camera Intrinsic and Extrinsic Parameters Calibration

Multi-camera sensor fusion requires an accurate camera projection modelling as well as accurate inter-camera extrinsic transforms. Given the wide Field-of-View and strong distortion of the fisheye camera lenses, we employ the equidistant distortion model (Kannala and Brandt 2006). Adjacent cameras in our setup share overlapping view frustums (approximately 36° horizontally), enabling the extrinsic calibration between the cameras via co-detection of known calibration target features. We calibrate both the

[†]https://github.com/sevensense-robotics/core_research_manual/




Figure 4. LiDAR point clouds overlaid on the camera images. This demonstrates the quality of camera intrinsics calibration and camera-LiDAR extrinsics calibration. In the left camera, the regions of the building without LiDAR points are due to the LiDAR's limited sensing range. Note that the motion undistortion process produces the jagged discontinuity in the LiDAR beam pattern in the right camera.

intrinsic and extrinsic parameters of the three cameras using the Kalibr open-source camera calibration toolbox (Rehder et al. 2016). The resultant calibration achieves sub-pixel reprojection error, as summarised in Tab. 2. We provide the camera calibration sequences with this dataset to facilitate experimentation with alternative calibration methods.

4.2 IMU Calibration

To facilitate visual and lidar odometry using the IMU, it is crucial to appropriately model the noise parameters of the IMU. For this, we measured the Allan variance parameters[‡] of the IMU accelerometer and gyroscope using an eight-hour data sequence.

4.3 Camera-IMU Extrinsic Calibration

With these accurate camera intrinsic parameters and IMU noise process parameters, we then perform camera-to-IMU extrinsic calibration individually for each of the cameras using Kalibr (Rehder et al. 2016). We cross-validated the consistency of the camera-IMU calibration by measuring the variation in the estimated coordinates of the IMU, using the individual camera-IMU extrinsic parameters and the camera-camera extrinsic parameters.

4.4 Camera-LiDAR Extrinsic Calibration

Camera-LiDAR extrinsic parameters are calibrated in a bundled fashion with the inter-camera extrinsic parameters from 4.1 held constant; a single $SE(3)$ transform between the bundle of cameras and the LiDAR is calibrated. We perform this calibration using DiffCal (Fu et al. 2023). This method uses a differentiable representation of the checker pattern to align the point intensities observed by the LiDAR directly with the camera-detected checkerboard pattern. In Fig. 4, We present an example of the LiDAR point




Figure 5. This image shows LiDAR overlay on top of the right-facing camera (cam2) demonstrating the effects of accurate LiDAR ego-motion undistortion. The motion here is approximately 0.9 m/s and 30 °/s, and the scene depth ranges between 1 m and 10 m. The right camera overlaps the start and end of the LiDAR sweep (the seam shown in green). The top figure shows the unprocessed raw point cloud overlay. The LiDAR points are slightly misaligned closer to the start of the sweep (b), and significantly misaligned at the end of the LiDAR sweep (a). With IMU-only point cloud undistortion, we achieve consistent overlays (bottom figure) (c) and (d), demonstrating the accuracy of the multi-modal spatiotemporal calibration.

clouds overlaid on the camera images using the described calibration. Furthermore, Fig. 5 illustrates the accuracy of the IMU-aided LiDAR point cloud motion undistortion by comparing the projection of raw and motion-undistorted LiDAR points onto the right-facing camera.

Calibration	Mean (px)	Median (px)	σ (px)
Camera/Camera			
cam0 (front)	0.22	0.21	0.14
cam1 (left)	0.23	0.22	0.15
cam2 (right)	0.22	0.21	0.14
Camera/IMU			
cam0 (front)	0.25	0.23	0.15

Table 2. Mean, median, and standard deviation of reprojection residuals for different calibration types.

5 Dataset

5.1 Data Format

The Oxford Spires Dataset consists of data collected in six sites in Oxford, UK, with multiple sequences taken at each

[‡]https://github.com/ori-drs/allan_variance_ros.




Figure 6. File structure of the Oxford Spires dataset: For *each sequence*, we provide the raw images and LiDAR point clouds, ground truth trajectory, LiDAR SLAM trajectory (including undistorted point clouds synchronised to images), and COLMAP trajectory. For *each site*, we provide TLS clouds as reconstruction ground truth. The ground truth systems are highlighted in red. We also provide calibration files for the camera intrinsics.

site (Sec. 5.2). The data is originally collected as rosbags[§]. We also provide raw sensor data (as individual files) as well as processed data. The processed data includes outputs from an example LiDAR SLAM system and a SfM system. Finally, we also provide the ground truth trajectories and reconstruction.

The following sections describe the raw data formats and the folder system which we provide for easy use of the data outside of ROS (Fig. 6).

5.1.1 Raw - Camera Images: The 20 Hz raw colour fisheye image streams from the three cameras of the Frontier are debayered and stored as 8-bit JPEG images. The three cameras are hardware-synchronised with each other, and hence the image triplets have the same timestamps. The images are stored as `<time>.jpg` under each camera folder, namely `cam0`, `cam1` and `cam2`, which correspond to the camera facing forward, left and right respectively. We debayer the raw image with bilinear interpolation, and white-balance the debayered image using the convolutional colour constancy method (Barron and Tsai 2017). We provide tools to white-balance[¶] the debayered images.

5.1.2 Raw - 3D LiDAR Point Clouds: 3D point clouds were collected using a Hesai QT64 LiDAR at 10Hz, and stored as `<time>.pcd`. Note that the unprocessed point clouds are raw measurements from the continuously scanning LiDAR, and the timestamp is the start time for each sweep. A subset of the LiDAR point clouds that were output by an example LiDAR SLAM system are also provided and described in Sec. 5.1.4.

5.1.3 Raw - IMU Measurements: The linear acceleration and angular velocity measurements from the IMU are

stored in `imu.csv`. Each row has the format of timestamp (seconds, nano-seconds), acceleration (x,y,z) and angular-velocity (x,y,z).

5.1.4 Processed - VILENS-SLAM Outputs: We provide the estimated trajectory and the motion undistorted point clouds output by LiDAR-inertial SLAM (VILENS-SLAM (Wisth et al. 2023; Ramezani et al. 2020a)). The trajectory is saved as `slam-poses.csv` with a $SE(3)$ pose estimate consisting of position (x,y,z) and quaternion (x,y,z,w) for each timestamp (seconds, nano-seconds). This data format can be directly used to test 3D reconstruction.

5.1.5 Processed - COLMAP Outputs: A solution to SfM is required as input to both MVS methods and radiance field methods (NeRF and 3D Gaussian Splatting). To facilitate researchers, we ran the state-of-the-art SfM method COLMAP (Schönberger and Frahm 2016) for each sequence and provide its outputs. Specifically, COLMAP provides camera information in `cameras.bin`, image information in `images.bin`, 3D feature points in `points3D.bin` and the database information in `database.db`.

Running SfM for all images captured at 20 Hz results in a large amount of output data and computation time. This is unnecessary because consecutive images are very close to each other and thus redundant. To keep the number of images manageable yet providing enough viewpoints for visual reconstruction, we selected images that are synchronised to the SLAM pose graph point clouds and spaced 1 m apart. We

[§]We provide rosbags in ROS1 and ROS2 format

[¶]Tool based on https://github.com/leggedrobotics/raw_image_pipeline




Figure 7. *Top:* Image rendered from TLS map in Keble College. *Bottom:* Depth image corresponding to the rendered image.

then ran COLMAP on this set of images. At walking speed, this results in a frequency of about 1 Hz. For each sequence, the total number of images was less than 2000 (for each of the three cameras). Using images aligned with corresponding LiDAR depth is also useful for methods that fuse LiDAR and vision, for example, colourising point clouds and depth-aided radiance fields such as Urban Radiance Field (Rematas et al. 2022). These aligned depth images are also provided in addition to the COLMAP outputs.

For compatibility with Nerfstudio (Tancik et al. 2023) (a popular open-sourced code base for state-of-the-art radiance field methods), we also convert the outputs from COLMAP into a `transforms.json` file. Specifically, `transforms.json` includes camera parameters (camera models, focal length, principle point, image size, distortion parameters), image file path and the corresponding $SE(3)$ pose estimate as a 4×4 transformation matrix.

Correcting the metric scale of vision-based 3D reconstructions produced by MVS and radiance field methods is necessary to enable comparison to the metric ground truth. To estimate the scale, we used Umeyama’s method^{||} to estimate a $Sim(3)$ transformation between the LiDAR trajectory and a COLMAP trajectory, and the results are saved in `evo.align.results.json`. We provide tools to compute the scale parameter and to rescale the MVS reconstruction and radiance field reconstruction to metric size.

5.1.6 Ground Truth - Reconstruction: We provide the registered individual TLS scans from Leica RTC360 for each site as the ground truth reconstruction. Each scan is saved as `<site-00x>.e57` under each site folder, and contains

not only the point clouds but also the sensor origin, which is important in reconstruction methods such as occupancy mapping (Hornung et al. 2013). Moreover, we also provide the complete colourised TLS map at 1 cm resolution for each site (Fig. 1) by merging the individual RTC360 scans. The TLS scan registration reports from the Leica software are provided as `FinalizeReport.pdf` in each ground truth site folder, which specifies the number of scans, overlap, residual in the merging process, etc. Regarding the density, we tried to distribute the scans uniformly on each site.

5.1.7 Ground Truth - Localisation: The ground truth trajectory is computed by ICP registering each undistorted LiDAR point cloud (as described in Sec. 5.1.4) to the merged TLS map described in Sec. 5.1.6. To obtain the ground truth pose for each LiDAR scan, we use an offline version of VILENS (Wisth et al. 2023) with Iterative Closest Point (ICP) (Besl and McKay 1992) at its core. Running offline, we can allocate enough time to register the LiDAR to the colourised TLS map. The trajectory estimated through the described process is synchronised with cameras by the procedure described in Sec. 3.1. This is in the same manner as for Newer College (Ramezani et al. 2020b) and Hilti-2022 (Zhang et al. 2022). The accuracy of the ground truth trajectory is approximately 1-2 cm. We validated the ground truth trajectory by projecting the individual LiDAR scans into a map and comparing them to the TLS map. The trajectory is provided as `gt-tum.txt` in TUM (Sturm et al. 2012) format, with each line encoding timestamp, position (x,y,z) and quaternion (x,y,z,w).

5.1.8 Depth images from the TLS map: Additionally, we include ground truth depth images rendered from the TLS map using the ground truth sensor trajectories. In Fig. 7, we show an example of an image rendered from the Keble College site with its corresponding depth image. These images could be used for evaluating monocular depth estimation or novel view rendering.

5.2 Sequence Description

The dataset was recorded in six historic sites in Oxford, UK:

- Bodleian Library ($\sim 37\,000\text{ m}^2$)
- Blenheim Palace ($\sim 14\,000\text{ m}^2$)
- Christ Church College ($\sim 26\,000\text{ m}^2$)
- Keble College ($\sim 18\,000\text{ m}^2$)
- Radcliffe Observatory Quarter (ROQ) ($\sim 12\,000\text{ m}^2$)
- New College ($\sim 18\,000\text{ m}^2$)

Each sequence was collected by walking with the Frontier payload device mounted in a backpack as shown in Fig. 3.

5.2.1 Bodleian Library: This site consists of the area around the Bodleian Library, which includes Radcliffe Square, where the Radcliffe Camera and Oxford’s University Church are located. It also reconstructs the outside of the Sheldonian Theatre and some of Broad Street. This part of the dataset contains the most iconic landmarks of the historic centre of Oxford (Fig. 1).

^{||}We used implementation from <https://github.com/MichaelGrupp/evo>




Figure 8. Examples of SLAM trajectories (in red) and LiDAR point cloud maps (in blue) for four sequences from the dataset.

For this site, we provide two outdoor trajectories of walking through streets and squares around the described area. The recordings contain many details of the predominant medieval and Gothic buildings.

5.2.2 Blenheim Palace: This is one of England’s largest houses and is notable as Sir Winston Churchill’s ancestral home. Five trajectories were captured in the palace’s main square, the principal hall, and rooms in the west wing, including the library. The trajectories have outdoor and indoor parts, including different-sized rooms and corridors. In Fig. 8 (b), we show an example of a sequence in the palace’s main square.

5.2.3 Christ Church College: Founded in 1546, Christ Church is a constituent college of Oxford and one of the city’s best-recognised locations. It contains Tom Quad, the largest square in Oxford, the college dining hall as well as Christ Church Cathedral and its cloister.

The sequences recorded in this site include outdoor areas with pavements and lawns as well as indoor parts with different lighting conditions and stairs accessing different levels, including the dining hall. One sequence is shown in Fig. 8 (a), which includes a complete loop of the perimeter of Tom Quad, which is challenging due to the limited range of the LiDAR sensor and repeating architecture.

5.2.4 Keble College: Keble is another constituent college of the University of Oxford. It comprises neo-Gothic-style buildings, including a hall and a church. Keble’s buildings are distinctive because they are constructed of alternating red and white coloured bricks — which provides an interesting challenge to visual reconstruction. This differentiates it from the other sites which are mostly built from limestone.

The Keble sequences were recorded outdoors in the college’s squares (Fig. 8 (c)), which include lawns and trees, as well as some interior and exterior parts.

5.2.5 Radcliffe Observatory Quarter: The ROQ site consists of the Faculty of Philosophy, the Mathematical Institute, and St Luke’s Chapel. This area is near the Oxford Robotics Institute, where the authors are affiliated. The two sequences recorded here contain squares with pavement, lawns, trees, narrow spaces between some buildings, and a fountain containing fine 3D details. In Fig. 8 (d), we show an example of a sequence through this site.

5.2.6 New College: New College is another constituent college of the University of Oxford and is located in the city’s historic centre. It contains squares, a hall, a church and a cloister. We recorded four sequences in New College, including an oval lawn area at the centre of the main quad surrounded by medieval buildings. The sequences combine outdoor and indoor parts with abrupt changes in light conditions. Most of the 03 sequence was walking through the park, containing a lawn and many trees. Some parts were fully covered by tree canopies. This site corresponds to the earlier New College Dataset (Smith et al. 2009; Ramezani et al. 2020b).

6 Benchmarks and Results

In this section, we describe three benchmarks we have created to demonstrate our dataset. The benchmarks compare state-of-the-art methods for localisation (Sec. 6.1), 3D reconstruction (Sec. 6.2), and novel view synthesis (Sec. 6.3).

6.1 Localisation Benchmark

In the localisation benchmark, we evaluate the trajectories estimated for each sequence using state-of-the-art LiDAR Inertial SLAM and SfM approaches:

- VILENS-SLAM: VILENS (Wisth et al. 2023) with pose graph optimisation (Ramezani et al. 2020a) (online).
- Fast-LIO-SLAM (Kim et al. 2022): Fast-LIO2 (Xu et al. 2022) with pose graph optimisation and Scan Context loop closures (Kim and Kim 2018) (online).
- SC-LIO-SAM (Kim et al. 2022): LIO-SAM (Shan et al. 2020) with pose graph optimisation and Scan Context loop closures (Kim and Kim 2018) (online).
- ImMesh (Lin et al. 2023): LiDAR meshing with Fast-LIO2 odometry (online).
- Fast-LIVO2 (Zheng et al. 2024): LiDAR visual inertial odometry (online).
- HBA (Liu et al. 2023): LiDAR bundle adjustment using as input the VILENS-SLAM result (offline).
- COLMAP (Schönbberger and Frahm 2016): Structure-from-Motion using only images (offline). For image matching, we used a sequential matcher with loop closure detection.

Table 3. Results in RMS (m) of the ATE using the provided ground truth. We mark the best results in blue using different tints. SC-LIO-SAM fails on some sequences. COLMAP gives incomplete results on some sequences.

Site	Sec	Len	VILENS-SLAM	Fast-LIO-SLAM	SC-LIO-SAM	ImMesh	Fast-LIVO2	HBA	COLMAP
Kemble College	02	290	0.06	0.25	1.26	0.08	0.95	0.11	0.05
	03	280	0.14	0.11	4.02	0.14	0.06	0.12	0.05
	04	780	0.16	0.49	✗	3.67	0.09	0.12	0.07
	05	710	0.11	0.29	✗	0.13	0.11	0.13	0.09
Radcliffe Obs. Quarter	01	400	0.06	0.17	0.23	0.20	0.04	0.05	0.07
	02	390	0.09	0.24	0.14	0.27	0.07	0.08	0.08
Blenheim Palace	01	490	0.47	0.18	6.74	0.27	0.14	0.21	0.08
	02	390	0.16	0.12	4.41	0.36	0.22	0.08	0.05
	05	390	1.05	0.28	✗	0.22	0.26	0.14	0.26
Christ Church College	01	920	0.06	0.72	✗	0.19	0.54	0.07	0.06
	02	640	0.17	0.49	✗	1.70	0.63	0.12	0.15
	03	340	0.03	0.23	0.14	0.16	✗	0.05	0.07
	05	820	0.17	0.30	✗	0.21	0.15	0.12	✗
Bodleian Library	02	690	1.11	0.25	1.71	0.39	0.46	0.89	0.27

We evaluate sequences on the trajectories which were completely within the ground truth TLS map. We exclude Kembel College 01, Blenheim Palace 03 and 04, Christ Church College 04 and 06, and Bodleian Library 01 from this analysis.




Figure 9. A top-down view showing a representative performance of the different systems for Sequence 01 at Blenheim Palace. The sequence starts and ends in the lower left. The environment where this sequence was collected can be seen in Fig. 8 (b).

6.1.1 Evaluation Metrics: In the dataset tools, we provide a Python script to compute Absolute Trajectory Error (ATE) and Relative Pose Error (RPE) metrics. We used ATE to compare the poses estimated by the SLAM and SfM methods as they should be globally consistent (Tab. 3). Even though Fast-LIVO2 is an odometry system, its drift rate is low. Thus, we consider it globally consistent and also evaluated it with ATE.

To transform the trajectories estimated by the methods to the ground truth coordinate frame (Sec. 5.1.7), we use the $SE(3)$ Umeyama alignment. For comparison between odometry systems, we recommend using RPE to measure local performance.

6.1.2 Experimental Results: A comparison between the listed methods using the provided ground truth (Sec. 5.1.7) is presented in Tab. 3 for each sequence. We mark the best results in dark blue, while the second and third are highlighted with a lighter tint. The offline methods on the right side (HBA and COLMAP) can take advantage of all available data and provide the most accurate results. While one might expect that HBA should provide better accuracy (as a LiDAR bundle adjustment approach), we achieved the best performance with COLMAP. We believe this is due to the overlapping multicamera configuration, which provides abundant view constraints and allows features to be seen from different perspectives with proper light conditions in general. Furthermore, the prior map we provide to HBA is imperfect which means that the undistorted point clouds can be imperfect, which would then lead in turn to inaccuracies in HBA. However, note also that COLMAP was unable to solve some sequences (with poor lighting) at all and instead produced multiple disconnected sub-models. This is usually due to insufficient visual features being matched in the area where two sub-models ought to connect, which can be due to there being insufficient visual overlap or there being challenging lighting conditions (e.g., The dining hall in Christ Church College is relatively dark). In comparison, the LiDAR SLAM systems are invariant to lighting conditions and the present visual features.

The second most accurate method is HBA, which further refines VILENS-SLAM’s trajectory estimation with LiDAR bundle adjustment in post-processing. Of the online methods, VILENS-SLAM and Fast-LIVO2 performed best. Fast-LIVO2 performs better in short-loop sequences (Kemble College and Radcliff Observatory Quarter), as the map created is also seen, and its drift is low. In contrast, VILENS-SLAM performs better in large loops as it performs loop closures. ImMesh and Fast-LIO-SLAM use Fast-LIVO2 (Xu et al. 2022) as their core odometry module, which achieved accurate trajectory estimation. SC-Fast-LIO performs loop

closure correction using Scan Context (Kim and Kim 2018), while ImMesh can recover from drifts while generating the mesh map. SC-LIO-SAM produces satisfactory results on some sequences using Scan Context (Kim and Kim 2018) as an appearance-based place recognition module. However, it adds incorrect loop closures in sequences with large loops and repeated building patterns, such as Christ Church College and Blenheim Palace. We note that these methods could potentially perform better with further parameter tuning.

In Fig. 9, we show a representative example of the performance of the evaluated methods using Sequence 01 of Blenheim Palace. All of the methods produce an accurate trajectory except for SC-LIO-SAM, which incorporates an incorrect loop closure when closing the large loop.

6.2 3D Reconstruction Benchmark

The reconstruction benchmark evaluates outputs from the systems that use vision or LiDAR. Specifically, we evaluate the following systems:

- VILENS-SLAM: Merged LiDAR point clouds using poses obtained using the LiDAR SLAM system described in Sec. 6.1 (online).
- OpenMVS^{**}: an MVS system which uses input from COLMAP (Schönberger and Frahm 2016) (offline).
- Nerfacto: The default and recommended method from Nerfstudio^{††} (Tancik et al. 2023) that combines features from MipNeRF-360 (Barron et al. 2022), Instant-NGP (Müller et al. 2022) and others. It uses input from COLMAP, as with OpenMVS (offline).

The outputs from each system are all in the form of 3D point clouds. For Nerfacto, the point cloud is generated from the trained model by calculating the expected depth and colour for the training rays, and projecting the depth points into 3D.

We selected example trajectories that are completely within the ground truth reconstruction from Blenheim Palace, Christ Church College, Keble College and Radcliffe Observatory Quarter.

6.2.1 Evaluation Metrics: We use the F-score as the primary metric for reconstruction. The F-score is calculated as the harmonic mean of precision and recall, thus it considers both aspects of the reconstruction: accuracy and completeness. To calculate precision and recall, we consider a point to be a true positive (TP) if the distance from it to the closest ground truth point is within a certain threshold. We report results using 5 cm and 10 cm thresholds. False positives (FP) are reconstructions that are further from the ground truth and thus inaccurate. False negatives (FN) are regions in the ground truth that have no neighbouring points in the reconstruction, and are thus incomplete. Specifically, precision and recall are defined by

$$\text{Precision} = \frac{TP}{TP + FP} \quad (1)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2)$$

The F-score is then calculated as

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

We also report the point-to-point distances to measure accuracy and completeness following the conventions of the DTU dataset (Aanæs et al. 2016). Specifically, accuracy is measured as the point-to-point distance from the reconstruction to the ground truth map and measures the reconstruction quality (the lower, the better). Completeness is the distance from the point-wise reference map to the reconstruction and indicates how much of the ground truth surface has been captured by the reconstruction (the lower, the better). However, these metrics are more sensitive to outliers and non-overlapping regions, which may skew the results in practice. Thus, we do not use them as the primary metrics.

6.2.2 Reconstruction Filtering: In practice, the reconstruction and ground truth reference models will contain regions which were not mutually scanned. If unaccounted for, this would lead to erroneous false positives and false negatives. In turn, this would result in precision and recall measures which do not reflect the true quality of the reconstruction. For fairer comparison, we filter out points in the reconstruction that fall outside the reconstructed ground truth region, i.e. the regions not reconstructed in the ground truth model. In particular, for Nerfacto the sky must be specifically removed because, as a dense representation, it attempts to reconstruct it using available depth cues. We filter these sky point clouds to ensure the evaluation focuses on the reconstruction of the physical environment itself.

6.2.3 Experimental Results: The quantitative evaluations of the reconstructions are presented in Tab. 4, and the qualitative results are shown in Fig. 10. VILENS-SLAM reconstruction achieves the best F-score in most experiments except Keble-04. This is a reasonable result given LiDAR’s accurate depth measurements. The remaining inaccuracy mostly comes from trajectory errors in long sequences and the presence of dynamic objects. The reconstruction completeness is limited by the short sensor range in Christ Church College and Blenheim where the central region of the large squares is not reconstructed.

Reconstructions from OpenMVS are accurate in regions with abundant view constraints and distinct texture, but it is not able to reconstruct surfaces with uniform texture such as the ground in Blenheim Palace and the lawn in Christ Church College. The error distribution in the MVS cloud is not uniform and tends to appear at surface boundaries where occlusion is an issue.


Although both OpenMVS and Nerfacto are purely vision-based reconstruction methods, Nerfacto point clouds are generally less precise. This is because MVS filters uncertain points (by checking photo-consistency), but the NeRF approach instead optimises a continuous radiance field without an explicit notion of uncertainty. For regions with insufficient view constraints and uniform texture, Nerfacto estimates incorrect depth values which leads to uneven

^{**}<https://cdcseacave.github.io/openMVS>

^{††}We used version 1.1.4

Table 4. Quantitative evaluation of the 3D reconstructions from VILENS-SLAM, OpenMVS and Nerfacto. We indicate the best results with a dark blue background.

Site	SEC.	Method	Accuracy↓	Completeness↓	Precision	5cm Recall	F-score	Precision	10cm Recall	F-score
Blenheim Palace	05	VILENS-SLAM	0.070	0.506	0.670	0.392	0.495	0.867	0.661	0.750
		OpenMVS	0.126	1.045	0.451	0.251	0.323	0.574	0.381	0.458
		Nerfacto	0.302	0.676	0.232	0.094	0.134	0.388	0.257	0.309
Christ Church College	02	VILENS-SLAM	0.082	3.296	0.540	0.250	0.342	0.794	0.408	0.539
		OpenMVS	0.046	5.381	0.771	0.201	0.319	0.886	0.266	0.410
		Nerfacto	0.219	4.435	0.328	0.157	0.212	0.532	0.254	0.343
Kemble College	04	VILENS-SLAM	0.067	0.342	0.527	0.527	0.527	0.816	0.779	0.797
		OpenMVS	0.050	0.409	0.766	0.606	0.677	0.918	0.718	0.806
		Nerfacto	0.137	0.150	0.418	0.484	0.449	0.654	0.709	0.680
Radcliffe Obs. Quarter	01	VILENS-SLAM	0.047	0.233	0.708	0.536	0.610	0.909	0.806	0.854
		OpenMVS	0.048	0.622	0.745	0.470	0.577	0.902	0.618	0.734
		Nerfacto	0.197	0.398	0.415	0.395	0.405	0.587	0.598	0.592

**Figure 10.** Comparison between the reconstructions achieved by the different methods. The reconstructions in the first three columns are coloured by point-to-point distance to the ground truth model.

ground reconstructions. In comparison, OpenMVS filters some of the reconstruction there, which leads to better precision and accuracy.

The reconstruction quality is determined not only by the reconstruction method but also by the accuracy of the input trajectory. Both precision and recall can be affected by an imperfect trajectory estimation. Clouds produced by VILENS-SLAM contain surfaces with high error that are the result of incorrectly registered LiDAR scans. Meanwhile, for Christ Church College, both the OpenMVS and Nerfacto reconstructions do not contain the dining hall (bottom left in the corresponding reconstructions from Fig. 10). This is because the dining hall could not be registered with the outdoor square by COLMAP (partly due to the poor lighting conditions as explained in 6.1.2).

6.3 Novel View Synthesis

We evaluate the quality of novel-view synthesis using the radiance field methods. Specifically, we evaluate:

- Nerfacto (Tancik et al. 2023) which is described in Sec. 6.2.
- Splatfacto (Ye et al. 2024), an implementation of 3D Gaussian Splatting (Kerbl et al. 2023) with quality comparable to the original implementation.

We also include results using the above methods with increased representation capability, namely Nerfacto-big (Nerfacto with larger hash grid size and proposal network size, and more ray samples) and Splatfacto-big (Splatfacto with lower thresholds for densifying and culling 3D Gaussians, which results in more Gaussians being used). All methods are trained for 5000 iterations. We select one in every ten images as the in-sequence evaluation images.

Table 5. Quantitative evaluation of Novel View Synthesis. The best results are coloured in blue using different tints. The test images are selected from the input trajectory (In-Sequence) as well as a separate trajectory with viewpoints far from the input trajectory (Out-of-Sequence).

Sequence	Method	In-Sequence			Out-of-Sequence		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Observatory Quarter	Nerfacto	23.40	0.807	0.336	21.25	0.786	0.370
	Nerfacto-big	20.66	0.807	0.292	19.38	0.787	0.317
	Splatfacto	22.76	0.791	0.373	19.47	0.736	0.445
	Splatfacto-big	23.54	0.811	0.347	20.26	0.761	0.413
Blenheim Palace	Nerfacto	18.42	0.716	0.506	17.09	0.682	0.537
	Nerfacto-big	17.93	0.724	0.445	17.09	0.695	0.493
	Splatfacto	19.34	0.726	0.589	16.02	0.668	0.659
	Splatfacto-big	19.77	0.733	0.576	16.20	0.671	0.643
Keble College	Nerfacto	21.10	0.731	0.397	20.29	0.748	0.368
	Nerfacto-big	19.71	0.749	0.326	18.15	0.736	0.381
	Splatfacto	20.47	0.651	0.514	19.92	0.658	0.500
	Splatfacto-big	21.36	0.688	0.478	20.86	0.707	0.434




Figure 11. Illustrative results of Splatfacto-big when evaluated using in-sequence (green) and out-of-sequence (red) trajectories. When the rendering viewpoint is quite different from the training trajectory, the rendered images exhibit many more artefacts. The in-sequence and out-of-sequence trajectories in Radcliffe Observatory Quarter and Blenheim Palace are in different directions, while the trajectories in Keble College have similar viewing directions but are from distant positions. From our test, we found that Splatfacto-big generates more visual artefacts than Nerfacto-big.

6.3.1 Evaluation Metrics: We measure the quality of the rendered images using the Peak Signal-to-noise Ratio (PSNR), Structural Similarity (SSIM) (Wang et al. 2004) and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al. 2018) metrics, as commonly used in the literature (Mildenhall et al. 2021; Barron et al. 2022; Tancik et al. 2023).

6.3.2 Out-of-Sequence Novel View Synthesis: When evaluating radiance field methods, methods often use test

poses that are close to the training poses. This is typically because the test poses and training poses are sampled from a common input trajectory. In downstream applications, the ability to render photorealistic images from viewpoints that are quite different from the training poses is crucial. To facilitate research in this direction, we generate challenging test sets whose viewpoints are very different from the training sets. Specifically, we merged images from different sequences taken in the same site using COLMAP. Then, we manually selected training and test set images that are

far away apart or have very different view directions. We describe the images that are selected from the input trajectory as “in-sequence” and images from a separate trajectory with different viewpoints as “out-of-sequence”.

6.3.3 Experimental Results: We present quantitative results in Tab. 5. Of particular interest, one can see that the quality of novel view synthesis falls significantly when moving from the in-sequence trajectory to the out-of-sequence trajectory. Compared to Nerfacto, Splatfacto (and its big version) generalises worse in the out-of-sequence setting, and we show qualitative results of Splatfacto-big in Fig. 11. The generalisation issue is particularly evident in Radcliffe Observatory Quarter and Keble College where the renderings are almost photorealistic from an angle close to the training data, but exhibit severe artefacts when rendered from a different location. Some of the artefacts have the wrong 3D geometry, and a typical issue is there being elongated 3D Gaussians along the training view angles as mentioned in Matsuki et al. (2024a). Other artefacts, such as the black artefact on the ground from Radcliffe Observatory Quarter, are due to the modelling of view-dependent colour used in radiance field methods. View-dependent colour is commonly modelled by a neural network (Mildenhall et al. 2021) or spherical harmonics (Kerbl et al. 2023). When the training viewing angles are limited (which is common in robotics applications), the optimised neural network or spherical harmonics can be overfitted which leads to unexpected colours when rendering from a novel viewing angle. This is a limitation of the state-of-the-art radiance fields method which is under-explored in the literature.

For the methods we tested, we found them all to be capable of generating reasonably photo-realistic images when rendering from in-sequence poses. A key difference between Nerfacto and Splatfacto is the rendering speed at test time: Both Splatfacto and Splatfacto-big render at 3.5 Hz on average, while Nerfacto renders at 1.25 Hz and Nerfacto-big at 0.57 Hz. When using the “big” version for Nerfacto and Splatfacto, the rendering quality is generally better with LPIPS increased by 9.6% and SSIM by 2% on average. This improvement is not always reflected in the PSNR measure, because it is also affected by the per-frame appearance difference (e.g. lighting) (Martin-Brualla et al. 2021) as illustrated in Fig. 12. For this reason, we give more consideration to changes in LPIPS and SSIM. The appearance difference issue can be potentially addressed by techniques such as test-time appearance encoding optimisation (Martin-Brualla et al. 2021).

7 Limitations and Future Work

The AlphaSense cameras have limited dynamic range. To be able to capture both dark and bright environments, we used the auto-exposure function of the cameras during the data collection. However, images captured using auto-exposure will have inconsistent pixel intensity when observing the same 3D structure from different viewpoints. Because of this, merging colourised LiDAR point clouds would lead to a mixture of different colours in the reconstruction. We believe this is an important research question, and there are several promising directions to address this issue, including estimating image exposure time (e.g. R3LIVE++ (Lin and




Figure 12. Comparison between an evaluation image and a rendered image from Nerfacto (Tancik et al. 2023). The PSNR metric is affected not only by the visual scene, but also by the lighting difference. Nerfacto uses per-frame appearance encodings (Martin-Brualla et al. 2021) which are optimised during training. When rendering at test time, Nerfacto uses the averaged appearance encoding of the training images. In the novel-view synthesis evaluation, we give more consideration to LPIPS and SSIM since they are more invariant to lighting differences.

Zhang 2024)) and modelling image appearance as latent features (e.g. NeRF-W (Martin-Brualla et al. 2021), GS-W (Zhang et al. 2024)).

Our localisation and reconstruction benchmarks primarily evaluate classical SLAM systems. The benchmarks can be extended to include more recent learning-based SLAM systems such as DROID-SLAM (Teed and Deng 2021), Gaussian Splatting SLAM (Matsuki et al. 2024b), MAST3R-SLAM (Murai et al. 2025) and PIN-SLAM (Pan et al. 2024). Evaluating these methods could provide more insights into their performance in large-scale outdoor data and facilitate the development of new approaches.

8 Conclusions

We present a large-scale dataset with colour images and LiDAR scans paired with high-quality ground truth 3D models and sensor trajectories. We demonstrate that the dataset is suitable for evaluating a variety of tasks in robotics and computer vision including LiDAR SLAM, Structure-from-Motion, Multi-View Stereo, Neural Radiance Field and 3D Gaussian Splatting. The scale of the provided data sequences and the quality of the ground truth trajectory and reconstruction make it suitable for evaluating large-scale localisation and 3D reconstruction methods in an outdoor environment. In addition, the colour cameras used in our dataset make it suitable for evaluating radiance field approaches, and encourage the development of SLAM systems integrated with radiance field representations. In particular, we demonstrate that state-of-the-art radiance field methods require further development to be applicable in the robotics context, namely inaccurate 3D geometry and limited generalisation capability when tested with poses distant from the training sequence.

Acknowledgements

The authors would like to thank Tobit Flatscher for helping with the GPS sensor, Prof. Ayoung Kim, Matias Mattamala and Christina Kassab for discussions and proofreading, Haedam Oh and Jianeng Wang for helping with dataset collection and calibration, Dongjae Lee for helping with post-processing.

Declaration of Conflicting Interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

This project has been partly funded by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT)(No. RS-2024-00461409). Miguel Ángel Muñoz-Bañón is supported by the Valencian Community Government and the European Union through the CIBEST/2023/44 fellowship and PROMETEO/2021/075 project.

References

- Aanæs H, Jensen RR, Vogiatzis G, Tola E and Dahl AB (2016) Large-scale data for multiple-view stereopsis. *Intl. J. of Computer Vision* 120: 153–168.
- Barron JT, Mildenhall B, Verbin D, Srinivasan PP and Hedman P (2022) Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)* : 5470–5479.
- Barron JT and Tsai YT (2017) Fast fourier color constancy. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 886–894.
- Besl PJ and McKay ND (1992) Method for registration of 3-d shapes. In: *Sensor fusion IV: control paradigms and data structures*, volume 1611. Spie, pp. 586–606.
- Brizi L, Giacomini E, Di Giammarino L, Ferrari S, Salem O, De Rebotti L and Grisetti G (2024) VBR: A vision benchmark in Rome. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*.
- Burri M, Nikolic J, Gohl P, Schneider T, Rehder J, Omari S, Achtelik MW and Siegwart R (2016) The EuRoC micro aerial vehicle datasets. *Intl. J. of Robotics Research* 35(10): 1157–1163.
- Carlevaris-Bianco N, Ushani AK and Eustice RM (2016) University of Michigan North Campus long-term vision and lidar dataset. *Intl. J. of Robotics Research* 35(9): 1023–1035.
- Chang A, Dai A, Funkhouser T, Halber M, Niessner M, Savva M, Song S, Zeng A and Zhang Y (2017) Matterport3D: Learning from RGB-D data in indoor environments. In: *IEEE Intl. Conf. on 3D Vision*. pp. 667–676.
- Dai A, Chang AX, Savva M, Halber M, Funkhouser T and Nießner M (2017) ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 5828–5839.
- Fu LFT, Chebrolu N and Fallon M (2023) Extrinsic calibration of camera to LIDAR using a differentiable checkerboard model. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. pp. 1825–1831.
- Geiger A, Lenz P, Stiller C and Urtasun R (2013) Vision meets robotics: The KITTI dataset. *Intl. J. of Robotics Research* 32(11): 1231–1237.
- Handa A, Whelan T, McDonald J and Davison AJ (2014) A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, pp. 1524–1531.
- Helmberger M, Morin K, Berner B, Kumar N, Cioffi G and Scaramuzza D (2022) The Hilti SLAM challenge dataset. *IEEE Robotics and Automation Letters* 7(3): 7518–7525.
- Hornung A, Wurm KM, Bennewitz M, Stachniss C and Burgard W (2013) OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots* 34: 189–206.
- Jeong J, Cho Y, Shin YS, Roh H and Kim A (2019) Complex urban dataset with multi-level sensors from highly diverse urban environments. *Intl. J. of Robotics Research* 38(6): 642–657.
- Jiang P, Osteen P, Wigness M and Saripalli S (2021) Rellis-3D dataset: Data, benchmarks and analysis. In: *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, pp. 1110–1116.
- Kannala J and Brandt S (2006) A generic camera model and calibration method for conventional, wide-angle, and fish-eye lenses. *IEEE Trans. Pattern Anal. Machine Intell.* 28(8): 1335–1340.
- Kerbl B, Kopanas G, Leimkühler T and Drettakis G (2023) 3D Gaussian Splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)* 42(4): 1–14.
- Kim G and Kim A (2018) Scan context: Egocentric spatial descriptor for place recognition within 3D point cloud map. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, pp. 4802–4809.
- Kim G, Yun S, Kim J and Kim A (2022) SC-LiDAR-SLAM: a front-end agnostic versatile LiDAR SLAM system. In: *International Conference on Electronics, Information, and Communication (ICEIC)*. IEEE, pp. 1–6.
- Knapitsch A, Park J, Zhou QY and Koltun V (2017) Tanks and Temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (TOG)* 36(4): 1–13.
- Li H, Zou Y, Chen N, Lin J, Liu X, Xu W, Zheng C, Li R, He D, Kong F, Cai Y, Liu Z, Zhou S, Xue K and Zhang F (2024) MARS-LVIG dataset: A multi-sensor aerial robots SLAM dataset for LiDAR-visual-inertial-GNSS fusion. *Intl. J. of Robotics Research* 43(8): 1114–1127.
- Li J, Wu W, Yang B, Zou X, Yang Y, Zhao X and Dong Z (2023) Whu-helmet: a helmet-based multisensor SLAM dataset for the evaluation of real-time 3D mapping in large-scale gnss-denied environments. *IEEE Transactions on Geoscience and Remote Sensing* 61: 1–16.
- Lin J, Yuan C, Cai Y, Li H, Ren Y, Zou Y, Hong X and Zhang F (2023) ImMesh: An immediate LiDAR localization and meshing framework. *IEEE Trans. Robotics* 39(6): 4312–4331.
- Lin J and Zhang F (2024) R³ live++: A robust, real-time, radiance reconstruction package with a tightly-coupled lidar-inertial-visual state estimator. *IEEE Trans. Pattern Anal. Machine Intell.* .
- Liu X, Liu Z, Kong F and Zhang F (2023) Large-scale LiDAR consistent mapping using hierarchical LiDAR bundle adjustment. *IEEE Robotics and Automation Letters* 8(3): 1523–1530.
- Liu Y, Fu Y, Qin M, Xu Y, Xu B, Chen F, Goossens B, Sun PZ, Yu H, Liu C, Chen L, Tao W and Zhao H (2024) BotanicGarden: A high-quality dataset for robot navigation in unstructured natural environments. *IEEE Robotics and Automation Letters* 9(3): 2798–2805.
- Martin-Brualla R, Radwan N, Sajjadi MSM, Barron JT, Dosovitskiy A and Duckworth D (2021) NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 7210–7219.

- Matsuki H, Murai R, Kelly PH and Davison AJ (2024a) Gaussian splatting SLAM. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 18039–18048.
- Matsuki H, Murai R, Kelly PH and Davison AJ (2024b) Gaussian splatting SLAM. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 18039–18048.
- Mildenhall B, Srinivasan PP, Ortiz-Cayon R, Kalantari NK, Ramamoorthi R, Ng R and Kar A (2019) Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* 38(4): 1–14.
- Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R and Ng R (2021) NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* 65(1): 99–106.
- Müller T, Evans A, Schied C and Keller A (2022) Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)* 41(4): 1–15.
- Murai R, Dexheimer E and Davison AJ (2025) MAST3R-SLAM: Real-time dense SLAM with 3D reconstruction priors. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*.
- Nair AD, Kindle J, Levchev P and Scaramuzza D (2024) Hilti SLAM Challenge 2023: Benchmarking single + multi-session SLAM across sensor constellations in construction. *IEEE Robotics and Automation Letters* 9(8): 7286–7293.
- Nguyen TM, Yuan S, Nguyen TH, Yin P, Cao H, Xie L, Wozniak M, Jensfelt P, Thiel M, Ziegenbein J and Blunder N (2024) MCD: Diverse large-scale multi-campus dataset for robot perception. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 22304–22313.
- Pan Y, Zhong X, Wiesmann L, Posewsky T, Behley J and Stachniss C (2024) PIN-SLAM: LiDAR SLAM using a point-based implicit neural representation for achieving global map consistency. *IEEE Trans. Robotics*.
- Ramezani M, Tinchev G, Iuganov E and Fallon M (2020a) Online LiDAR-SLAM for legged robots with robust registration and deep-learned loop closure. In: *IEEE Robotics and Automation Letters*. pp. 4158–4164.
- Ramezani M, Wang Y, Camurri M, Wisth D, Mattamala M and Fallon M (2020b) The Newer College dataset: Handheld lidar, inertial and vision with ground truth. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, pp. 4353–4360.
- Rehder J, Nikolic J, Schneider T, Hinzmann T and Siegwart R (2016) Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. pp. 4304–4311.
- Rematas K, Liu A, Srinivasan PP, Barron JT, Tagliasacchi A, Funkhouser T and Ferrari V (2022) Urban radiance fields. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 12932–12942.
- Sarlin PE, Dusmanu M, Schönberger JL, Speciale P, Gruber L, Larsson V, Miksik O and Pollefeys M (2022) LaMAR: Benchmarking localization and mapping for augmented reality. In: *Eur. Conf. on Computer Vision (ECCV)*. pp. 686–704.
- Schmidt F, Blessing C, Enzweiler M and Valada A (2024) ROVER: A multi-season dataset for visual SLAM. *arXiv preprint arXiv:2412.02506*.
- Schönberger JL and Frahm JM (2016) Structure-from-Motion revisited. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 4104–4113.
- Schops T, Schönberger JL, Galliani S, Sattler T, Schindler K, Pollefeys M and Geiger A (2017) A multi-view stereo benchmark with high-resolution images and multi-camera videos. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 3260–3269.
- Schubert D, Goll T, Demmel N, Usenko V, Stueckler J and Cremers D (2018) The TUM VI benchmark for evaluating visual-inertial odometry. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. pp. 1680–1687.
- Seitz SM, Curless B, Diebel J, Scharstein D and Szeliski R (2006) A comparison and evaluation of multi-view stereo reconstruction algorithms. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*, volume 1. IEEE, pp. 519–528.
- Shan T, Englot B, Meyers D, Wang W, Ratti C and Daniela R (2020) LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, pp. 5135–5142.
- Smith M, Baldwin I, Churchill W, Paul R and Newman P (2009) The New College vision and laser data set. *Intl. J. of Robotics Research* 28(5): 595–599.
- Straub J, Whelan T, Ma L, Chen Y, Wijmans E, Green S, Engel JJ, Mur-Artal R, Ren C, Verma S, Clarkson A, Yan M, Budge B, Yan Y, Pan X, Yon J, Zou Y, Leon K, Carter N, Briales J, Gillingham T, Mueggler E, Pesqueira L, Savva M, Batra D, Strasdat HM, Nardi RD, Goesele M, Lovegrove S and Newcombe R (2019) The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*.
- Sturm J, Engelhard N, Endres F, Burgard W and Cremers D (2012) A benchmark for the evaluation of RGB-D SLAM systems. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. pp. 573–580.
- Sucar E, Liu S, Ortiz J and Davison AJ (2021) iMAP: Implicit mapping and positioning in real-time. In: *Intl. Conf. on Computer Vision (ICCV)*. pp. 6229–6238.
- Sun T, Hao Y, Huang S, Savarese S, Schindler K, Pollefeys M and Armeni I (2023) Nothing Stands Still: A spatiotemporal benchmark on 3D point cloud registration under large geometric and temporal change. *arXiv preprint arXiv:2311.09346*.
- Tancik M, Weber E, Ng E, Li R, Yi B, Wang T, Kristoffersen A, Austin J, Salah K, Ahuja A, McAllister D, Kerr J and Kanazawa A (2023) Nerfstudio: A modular framework for neural radiance field development. In: *SIGGRAPH*. pp. 1–12.
- Teed Z and Deng J (2021) DROID-SLAM: Deep visual SLAM for monocular, stereo, and RGB-D cameras 34: 16558–16569.
- Wang Z, Bovik A, Sheikh H and Simoncelli E (2004) Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13(4): 600–612.
- Wei H, Jiao J, Hu X, Yu J, Xie X, Wu J, Zhu Y, Liu Y, Wang L and Liu M (2024) FusionPortableV2: A unified multi-sensor dataset for generalized SLAM across diverse platforms and scalable environments. *Intl. J. of Robotics Research*.
- Wisth D, Camurri M, Das S and Fallon M (2021) Unified multi-modal landmark tracking for tightly coupled lidar-visual-inertial odometry. *IEEE Robotics and Automation Letters* 6(2): 1004–1011.
- Wisth D, Camurri M and Fallon M (2023) VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots. *IEEE*

- Trans. Robotics* 39(1): 309–326.
- Xu W, Cai Y, He D, Lin J and Zhang F (2022) Fast-LIO2: Fast direct lidar-inertial odometry. *IEEE Trans. Robotics* 38(4): 2053–2073.
- Ye V, Li R, Kerr J, Turkulainen M, Yi B, Pan Z, Seiskari O, Ye J, Hu J, Tancik M and Kanazawa A (2024) gsplat: An open-source library for Gaussian splatting. *arXiv preprint arXiv:2409.06765*.
- Yeshwanth C, Liu YC, Nießner M and Dai A (2023) ScanNet++: A high-fidelity dataset of 3D indoor scenes. In: *Intl. Conf. on Computer Vision (ICCV)*. pp. 12–22.
- Yogamani S, Hughes C, Horgan J, Sistu G, Varley P, O’Dea D, Uricár M, Milz S, Simon M, Amende K et al. (2019) Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving. In: *Intl. Conf. on Computer Vision (ICCV)*. pp. 9308–9318.
- Zhang D, Wang C, Wang W, Li P, Qin M and Wang H (2024) Gaussian in the wild: 3D gaussian splatting for unconstrained image collections. In: *Eur. Conf. on Computer Vision (ECCV)*. Springer, pp. 341–359.
- Zhang L, Helmberger M, Fu LFT, Wisth D, Camurri M, Scaramuzza D and Fallon M (2022) Hilti-Oxford dataset: A millimeter-accurate benchmark for simultaneous localization and mapping. *IEEE Robotics and Automation Letters* 8(1): 408–415.
- Zhang R, Isola P, Efros AA, Shechtman E and Wang O (2018) The unreasonable effectiveness of deep features as a perceptual metric. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 586–595.
- Zheng C, Xu W, Zou Z, Hua T, Yuan C, He D, Zhou B, Liu Z, Lin J, Zhu F et al. (2024) Fast-LIVO2: Fast, direct lidar-inertial-visual odometry. *IEEE Transactions on Robotics*.
- Zhu Z, Peng S, Larsson V, Xu W, Bao H, Cui Z, Oswald MR and Pollefeys M (2022) NICE-SLAM: Neural implicit scalable encoding for SLAM. In: *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*. pp. 12786–12796.