Challenges of SLAM in extremely unstructured environments: the DLR Planetary Stereo, Solid-State LiDAR, Inertial Dataset

Riccardo Giubilato¹, Wolfgang Stürzl¹, Armin Wedler¹ and Rudolph Triebel^{1,2}

Abstract— We present the DLR Planetary Stereo, Solid-State LiDAR, Inertial (S3LI) dataset, recorded on Mt. Etna, Sicily, an environment analogous to the Moon and Mars, using a handheld sensor suite with attributes suitable for implementation on a space-like mobile rover. The environment is characterized by challenging conditions regarding both the visual and structural appearance: severe visual aliasing poses significant limitations to the ability of visual SLAM systems to perform place recognition, while the absence of outstanding structural details, joined with the limited Field-of-View of the utilized Solid-State LiDAR sensor, challenges traditional LiDAR SLAM for the task of pose estimation using point clouds alone. With this data, that covers more than 4 kilometers of travel on soft volcanic slopes, we aim to: 1) provide a tool to expose limitations of state-ofthe-art SLAM systems with respect to environments, which are not present in widely available datasets and 2) motivate the development of novel localization and mapping approaches, that rely efficiently on the complementary capabilities of the two sensors. The dataset is accessible at the following url: https://rmc.dlr.de/s3li_dataset

I. INTRODUCTION

Simultaneous Localization and Mapping (SLAM) techniques can nowadays be considered mature enough to find applications in many fields such as autonomous driving [1], [2], automated construction [3] or agriculture [4], [5]. While typically implemented with the use of multi-camera setups or RGB-D sensors, the recent democratization of 3D LiDAR sensors allowed easy integration of range sensing on mobile robots, enabling robust and large-scale mapping especially in urban or man-made scenarios.

The wide variety of published visual or LiDAR SLAM approaches [6]–[8] might suggest that the SLAM problem could be considered solved. However, we argue that this is far from the truth, specifically when dealing with environments that are unstructured and severely aliased [9], [10]. This is especially true in the case of planetary analogous environments, which are usually located in extreme locations such as deserts or volcanic surfaces. In this case, the lack of uniquely identifiable visual or structural features prevents from performing place recognition in a robust and repeatable manner. Moreover, when dealing with LiDAR SLAM, the lack of vertical structures leads to missing constraints for a robust convergence of ICP-style algorithms [11].

Typical datasets fail to address this extreme case. Among the most popular, KITTI [12], Oxford RobotCar [13], KAIST [14] and 4Season [15] tackle the case of autonomous driving



Fig. 1. Bird's-eye view of the site where the dataset is recorded with trajectories for each sequence overlaid on top. The Cisternazza crater is visible on the right. Tiles (C) Esri – Source: Esri, i-cubed, USDA, USGS, AEX, GeoEye, Getmapping, Aerogrid, IGN, IGP, UPR-EGP and the GIS User Community

in urban scenarios. While presenting sequences with significant challenges, e.g., dynamic objects or large time intervals spanning seasons, the point of view of a road vehicle, as well as the panoramic 3D LiDAR scans captured within highly structured scenarios, facilitate the tasks of computing odometry and performing place recognition.

Other datasets, such as TUM RGB-D [16], TUM VI [17] or the ETH3D SLAM benchmark [18] provide sequences recorded with hand-held stereo or RGB-D cameras principally in indoor environments. Their sequences tend to be short and frequently re-observing places from similar viewpoints and under very limited changes in the visual appearance.

To overcome the limitations related to the practical realization of the sensor setup, but also to model several motion characteristics and environments, many synthetic datasets have been recently released such as ICL-NUIM [19] or TartanAir [20]. While this has many advantages for the development of SLAM algorithms, e.g., the diversity of motion patterns and scenes is beneficial towards training of learningbased approaches, we believe that, especially for robots operating in the field, unstructured natural environments still offer the most challenging conditions for the implementation of visual or LiDAR-based SLAM.

Several datasets have been proposed which are recorded

¹German Aerospace Center (DLR), Institute of Robotics and Mechatronics, 82234 Weßling, Germany. firstname.lastname@dlr.de

²Department of Aerospace and Geodesy, Technical University of Munich, 80333 München, Germany

TABLE I SLAM datasets in unstructured environments

Ref	Sensors N seq L		Length	GT	Obs ⁴
[24]	Stereo LiDAR (3D)	3	$\sim 2 \text{ km}$ $\sim 100 \text{ min}$	D-GNSS	R
[22]	Stereo RADAR (2D)	43 ¹	$\sim 2 \ \mathrm{km}$	LiDAR ²	R
[29]	Stereo	2	$\sim 1 \text{ km}$	D-GNSS	R
[21]	Stereo	23^{3}	$\sim 10 \ {\rm km}$	D-GNSS	R
[28]	Stereo Omni	35	\sim 9.2 km	D-GNSS Heading	Н
Ours	Stereo LiDAR (Solid St.)	7	$\sim 4.3 \text{ km}$ $\sim 90 \text{ min}$	D-GNSS	Н

¹ Sequences with traverses, other 44 are static LiDAR scans

² High-res LiDAR scans captured statically at 30-40 meters distance

³ Splits of a long traverse without revisiting places

⁴ R: Rover, H: Handheld sensor assembly

from the perspective of a mobile robot operating in the field [21]. The Chilean Underground Mine dataset [22] provides sequences recorded in a mine with a stereo camera, a 3D LiDAR and a radar. The Visual-Inertial Canoe dataset [23] provides sequences where the main challenge is represented by reflections on the water. The Katwijk beach dataset [24] is recorded on a beach with artificial landmarks from a planetary-like rover equipped with a stereo camera that has similar characteristics to the one of the ExoMars rover. However challenging, the environment of caves and mines is, as was also shown during the DARPA SubT Challenge [25]-[27], particularly suited for LiDAR-based SLAM, as are the stone replicas on the flat sandy surface of the Katwijk beach. The MADMAX dataset [28], collected in the Morocco desert using a hand-held stereo unit with accurate ground truth, offers stereo sequences which are particularly challenging for traditional place recognition, but does not provide LiDAR data, which could be beneficial as it can provide information for localization invariant to the visual appearance. Similarly, [29] provides two sequences recorded on Mt. Etna with a mobile rover, which was only equipped with a stereo camera mounted on a pan-tilt unit.

We fill this gap and present here a dataset collected on the volcanic surface of Mt. Etna, Italy, a planetary analogous environment characterized by extreme visual aliasing and lack of distinct structural features. The dataset has been collected with a hand-held sensor assembly that includes a stereo monochrome camera setup, a Solid State LiDAR and an IMU. The setup replicates the typical height and point of view of the sensors mounted on a mobile robot. Although influenced by the walking pattern, the recordings look smooth and free of sharp or high-frequency motions. Seven sequences, 8 to 30 minutes long and covering distances up to 1.3 kilometers, have been recorded in order to evaluate the ability of SLAM algorithms:

- 1) to accurately estimate odometry under harsh lighting conditions and lack of unique features,
- 2) to rely on different sensing modalities, depending on the structure and appearance of the environment,
- 3) to use visual and structural information to robustly

perform place recognition under severe aliasing.

In summary, this paper presents the following contributions:

- We release a dataset to challenge SLAM algorithms in an natural, extremely unstructured environment, that resembles the conditions of planetary-like scenarios.
- We evaluate a selection of state-of-the-art visual, and visual-inertial SLAM algorithms.
- We provide examples of using the data in unconventional manners, leveraging the interaction of cameras and LiDARs towards more robust SLAM in severely unstructured scenes.

II. THE DATASET

The dataset was recorded on the volcanic environment of Mt. Etna, Sicily, an active stratovolcano, which is denoted as a planetary analogous site for Moon and Mars due to its peculiar geological features [30]. The dataset includes 7 sequences, that were recorded on a slope at 2650 meters above sea level and in proximity to the Cisternazza crater (see Fig. 1). Each sequence is located far from hiking trails and provides views of an extreme environment characterized by a surface of smooth dark lava ash, due to the frequent eruptions that happened overnight at the time of recording the data, extreme visual contrast, due to the darkness of the soil and the brightness of the sky, and scarcity of geological features that might facilitate localization with respect to prior observations of the same places. In comparison with other similar datasets, listed in Table I, we provide an honest multisensory look at an environment, analogous to potential sites for robotic planetary exploration, without alterations, e.g. by introducing artificial landmarks, and without a prior selection based on properties that help specific sensing modalities, i.e., flatness of the soil, presence of outstanding structural details, or presence of distinct visual textures on the ground.

A. Sequences

The 7 sequences provided can be categorized according to the challenges they pose for localization algorithms. s3li_traverse_1, s3li_traverse_2 and s3li_crater test the ability of localization methods or SLAM algorithms to provide accurate pose estimation during traverses. The first two observe mainly ash slopes, with the sporadic presence of rocks or boulders, while the third is a long traverse, that starts and ends at the initial location, around the rim of the crater, providing long range features for which depth is unknown. s3li_loops and s3li_crater_inout are, instead, recorded with the purpose of evaluating the ability of SLAM algorithms to perform place recognition, exploiting either the visual or structural similarity. While the first revisits mainly unstructured slopes, that might be recognized through carefully tuned visual loop closure detectors, the second includes short repeating traverses within the rim of the crater, that could be exploited with both sensing modalities. s3li_mapping instead tests the ability of SLAM algorithms to produce a complete and consistent map of the environment relying on stereo as well as LiDAR depth measurements, as a confined area



Fig. 2. Collection of shots from the 7 sequences of the dataset. For each sequence, two examples are shown that include an image from the left camera, the stereo disparity and the closest LiDAR scan in time. All images are shown after applying CLAHE (Contrast-Limited Adaptive Histogram Equalization), to enhance the contrast, and the disparity images are produced using OpenCV's SGBM with numDisparities = 96 and blockSize = 11.

is explored multiple times and there should be plenty of loop closure opportunities. Finally, s3li_landmarks allows testing of segmentation techniques for natural landmarks, with the aim of either semantic mapping or segment-based loop closure detection [31]. An overview of the sequences, including D-GNSS tracks, is provided in Table II.

B. The hand-held setup (S3LI)

The sensor setup used to collect the data, depicted in Fig. 3, is designed to be compliant to a potential implemen-

tation on a planetary-like rover [32] and comprises:

Stereo Camera: two AVT Mako cameras are configured in a stereo setup with baseline 20 centimeters. The cameras are triggered at a frequency of 30 Hz to capture monochromatic images with resolution of 688×512 pixels using automatic exposure control to cope with the harsh lighting conditions on the mountain. Their timestamps are synchronized to the system time using the PTP protocol, where cameras function as slave to the master clock of the laptop. Compared to other

TABLE IIOverview of the sequences

Name	Length	D-GPS Track ¹	Description
s3li_traverse_1	371 m 9 min	E -50 -100 -100 -100 -100 -100 -25 50 East [m]	 Traverse along a ridge of stones on a slight downhill. The environment is nicely contrasted, however the severe aliasing undermines the chances for visual place recognition. The path intersects briefly in the middle section but there are no evident loop closure opportunities, except at the very end. Challenging LiDAR clouds mainly due to the slope and the absence of distinct structures.
s3li_traverse_2	300 m 8 min	Ē 25 tu 0 50 100 150 East [m]	 Short traverse from the starting point to the rim of the Cisternazza crater and back. Several panning motions during the traverse help to cover more area for purposes of mapping. There are basically no loop closure opportunities. As for s3li_traverse_1, LiDAR clouds have limited extent and do not include many useful structures.
s3li_crater	1010 m 19 min	E 0 to 200 300 East [m]	 First long sequence around the rim of the crater, returning to the starting location. The visual appearance is diverse, as many different terrain types are observed including sandy slopes and the rocky walls within the crater rim. Opportunities for mapping the insides of the crater using the LiDAR, more predominant structures are present, although none can be used to close loops as they are not revisited.
s3li_loops	587 m 13 min	E 0 -50 -100 East [m]	 Loopy sequence around the rocky ridges observed also in s3li_traverse_1 and s3li_traverse_2, plenty of opportunities for visual loop closure detection and/or for visual or structure-based segmentation of natural features. Less monotonous visual appearance than s3li_traverse_1 and s3li_traverse_2, a few salient rocky features are observed in "orbiting" motions. Structure is less interesting than visual appearance.
s3li_crater_inout	1338 m 28 min	-200 -200 East [m]	 Long traverse around the crater, that also includes a short travel below the rim to observe some geological formations. They are revisited after completing the first traverse, offering opportunities for loop closure detection. Challenging visual appearance characterized by sharp changes in illumination due to the exploration of the insides of the crater, often lying in the shadows Interesting sequence to map geological formations with the LiDAR, to be exploited also towards structure-based place recognition.
s3li_mapping	242 m 8 min	$ \begin{array}{c} 20 \\ \underline{E} \\ 10 \\ \underline{H} \\ 0 \\ -10 \\ 0 \\ \underline{20} \\ East [m] \end{array} $	 Sequence of limited spatial extent targeted at mapping, where the sensors observe the environment to cover as much area as possible. The terrain is characterized by ash slopes with few small boulders. Plenty of loop closure opportunities are available, however the contrast between the dark lava ash and harsh lighting challenges traditional visual methods. LiDAR measurements cover a dense area and can be interesting for complementing the mapping ability of the stereo camera as well as for providing solid ground constraints in the context of multi-modal SLAM.
s3li_landmarks	482 m 9 min	E 50 E 50 0 100 East [m]	 Sequence of traverses between large and isolated rock formations, some of which are re-observed to allow for an easy recognition of structures and visual features. As for s3li_mapping, the visual appearance is challenging due to the harsh contrast between dark lava ash and the bright sky. During the traverses, the ground offers limited visual details. The rock formations are observed from many viewpoints allowing for a complete reconstruction. Salient structures can easily be segmented for semantic mapping or matching.

¹ The green color denotes D-GNSS estimates under RTK-FIX ($\varepsilon \sim 1$ cm), while the yellow color denotes RTK-FLOAT ($\varepsilon \sim \lambda_{carrier}$).



Fig. 3. The Stereo, Solid-State LiDAR, Inertial (S3LI) sensor setup

datasets based on stereo images [28], [29], we do not include depth images, as the computation of disparity is not intended as a sensor output, but as part of the utilization of the stereo images. Many SLAM architectures [33]–[35], in fact, although relying on a stereo image stream, do not make use of dense disparity images, but instead perform sparse feature matching when needed. Nevertheless, as an example and for the purpose of visualization, Fig. 2 presents a collection of disparity images obtained with OpenCV's SGBM stereo algorithm. Due to the automatic exposure control that, in combination with the extreme intensity differences between the dark soil and the bright sky, often compresses the ground details into dark patches, we recommend the usage of automatic contrast enhancement algorithms such as OpenCV's CLAHE.

LiDAR: we use a Blickfeld Cube-1 LiDAR, which employs a MEMS-actuated beam deflection mirror, instead of traditional 360° LiDARs, as this type of construction is more suited for potential space applications due to increased mechanical robustness and reduced weight and power consumption. For all sequences, the LiDAR is configured to capture point clouds with a maximum number of 17400 points with a Field-of-View of about 70°H \times 30°V, which results in a scan rate of 4.7 Hz. Each data point is timestamped with respect to the origin of the scan and contains additional intensity information that can be used, for instance, to identify low-quality measurements. The synchronization of the scan timestamps with respect to the time of the laptop is realized with NTP.

IMU: an XSens MTi-G 10, connected via USB, provides linear acceleration and angular velocity at a rate of 400 Hz. **GPS receiver**: an Ublox f9p GNSS receiver mounted on the hand-held setup provides accurate differential estimates in conjunction with a base station. The data logs are processed in a later stage using the RTKLIB to obtain ground truth positions at a frequency of 5 Hz and centimeter-level accuracy.

C. Sensor calibration

Stereo camera and IMU: The intrinsic and extrinsic parameters of the stereo camera pair are calibrated using Calde/Callab [37], while the IMU to camera calibration is performed with Kalibr [38]. In order to facilitate the playback of data, the stereo images are provided already rectified.



Fig. 4. Alignment of LiDAR and (left) camera reference frames. In the top row: examples of binary edge images and projected LiDAR scans after extrinsic calibration. The color scheme of the projected LiDAR points reflects the direction of the normal vectors. The other rows show examples of LiDAR scans, projected on the left camera image, respectively from sequences s3li_traverse_1, s3li_mapping and s3li_crater_inout. Point clouds are downsampled to avoid cluttering the images and color-coded according to the depth of points.

In fact, as the lenses do not have a particularly wide field of view, the images corrected for distortion do not lose a significant amount of visual detail.

LiDAR-Camera extrinsic calibration: the LiDAR is extrinsically calibrated with respect to the left camera. As the solid-state LiDAR that we employed is characterized by quite a large minimum measurable depth of about 4 meters, traditional techniques for LiDAR to camera calibration proved either non-effective, e.g., in the case of a

RMSE^{*} (normalized) and ratio of completion of the compared SLAM algorithms									
Algorithm	s3li_traverse_1	s3li_loops	s3li_crater	s3li_crater_inout	s3li_landmarks	s3li_mapping	s3li_traverse_2		
ORB-SLAM3 (S) [33]	0.86 (100)	0.21 (100)	0.87 (56.5)	0.19 (69.5)	-	0.71 (41.2)	0.38 (84.1)		
VINS_Fusion (S) [36]	1.41 (100)	6.23 (100)	2.25 (100)	-	3.25 (50.1)	2.01 (62.5)	1.33 (100)		
VINS_Fusion (SI) [36]	0.62 (100)	-	-	0.25 (100)	0.36 (100)	-	1.90 (100)		
OPEN_VINS (SI) [35]	0.54 (100)	0.41 (80.3)	0.54 (42.6)	0.16 (86.9)	0.77 (83.9)	-	0.39 (100)		
BASALT (S) [34]	0.75 (100)	0.35 (100)	2.81 (100)	1.08 (100)	1.34 (100)	0.47 (100)	0.57 (100)		
BASALT (SI) [34]	0.66 (100)	0.46 (100)	1.64 (100)	0.41 (100)	0.86 (100)	0.44 (100)	0.38 (100)		

TABLE III ${\sf RMSE}^* \ ({\sf normalized}) \ {\sf and} \ {\sf ratio} \ {\sf of} \ {\sf completion} \ {\sf of} \ {\sf the} \ {\sf compared} \ {\sf SLAM} \ {\sf algorithms}$

* In bold are highlighted the best results, from those algorithms that completed the sequence

checkerboard pattern printed on planar target [39], or too impractical to perform on the field, e.g., methodologies that rely on non-planar geometries [40], [41]. For this reason, the transformation between the camera and LiDAR reference frames is determined by fixing the translation component to the one obtained by a CAD model of the sensor assembly, while the rotation component is optimized in an automatic edge alignment scheme [42], [43]. More specifically, n pairs of left camera images and LiDAR scans are manually selected from recording of an indoor environment, characterized by structural regularity and scarcity of evident textures. From each image, strong Canny edges are selected to obtain binary images, where the identified edges are an input to the optimization scheme. From each corresponding LiDAR scan, we estimate the local curvatures, using the PCL library, and edge points, which lie in correspondence to depth discontinuities between neighboring points in the sequence of measurements of the LiDAR beam. Formally, the extrinsic calibration is solved as a non-linear least squares problem, where the goal is to find the optimal rotation between LiDAR and left camera, that aligns the projections of LiDAR points, lying on edges or discontinuities, to edges in the image:

$$\arg\min_{R_{\mathrm{L}}^{\mathrm{C}_{\mathrm{l}}}} \sum_{i=1}^{n} \sum_{k=1}^{m_{i}} \omega_{ik} \rho_{c} \big(\|\mathbf{x}_{ik} - \pi \big(R_{\mathrm{L}}^{\mathrm{C}_{\mathrm{l}}} (\mathbf{p}_{ik}^{\mathrm{L}} - \mathbf{t}_{\mathrm{C}_{\mathrm{l}}}^{\mathrm{L}}) \big) \| \big), \quad (1)$$

where *n* is the number of image-scan pairs, m_i is the number of correspondences for the *i*-th pair, π is the camera projection function, $\mathbf{p}_{ik}^{\text{L}}$ and \mathbf{x}_{ik} are, respectively, 3D LiDAR and 2D image points on edges, $R_{\text{L}}^{\text{C}_1}$ and $\mathbf{t}_{\text{C}_1}^{\text{L}}$ are the rotation (to be optimized) and translation (known form CAD model) of the LiDAR in the left camera reference frame, ρ_c is a Cauchy robust loss function, and ω_{ik} is a weight proportional to the local curvature of the point cloud. The association between LiDAR and image edge points is performed iteratively in a nearest-neighbor approach, and the problem is solved a few consecutive times reducing the value of parameter *c* for the Cauchy loss.

III. USE-CASES

A. Visual SLAM

A major application area of this dataset is to evaluate the performances of SLAM algorithms in environments with challenging visual appearance. While the texture of the terrain can ease the task of, for instance, sparse optical flow, feature matching and place recognition are particularly challenged by the severe visual aliasing. The algorithms compared here are ORB_SLAM3 [33], VINS_Fusion [36], OPEN_VINS [35] and BASALT [34] in stereo or stereoinertial configuration. The estimated path is correlated to the D-GNSS ground truth using time correspondences. After this step, we align all pairs of matching positions using Horn's algorithm, and compute the root mean square of all the residuals. To be comparable amongst all sequences, the RMSE is normalized by the length of the associated ground truth trajectory. Furthermore, as the RMSE is affected by the length of the estimated trajectory (shorter trajectories exhibit less translational and rotational drift) we provide, as an additional performance indicator, the proportion of each sequence for which poses are estimated from Visual SLAM. This value is computed as the length of the ground truth, associated to SLAM estimates, over its total length. Table II-C reports these metrics for all algorithms and for every sequence of the dataset. Compared to the others, ORB_SLAM3, where camera tracking depends heavily on ORB descriptors matching, exhibits very good accuracy in all datasets, but frequently fails to estimate the whole sequence, especially when the observed environment is characterized almost completely by a very repetitive terrain, without the presence of rocks or features of unique appearance. The other algorithms, where the localization front-end depends on tracking strong corners through optical flow, do not suffer from this specific problem, but are generally more prone to pose drift.

As expected, the lowest errors are scored from stereoinertial configurations, as the inertial measurements help significantly to deal with high accelerations and rotational velocities, typical of the walking motion patterns when recording with hand-held devices. However, compared to their stereo counterparts, the delicate procedures to initialize IMU parameters leads often to lose significant portions of the initial parts of the sequences. Furthermore, we experienced random behaviours of this step, especially with VINS_Fusion and OPEN_VINS, often initializing the state estimation at very different points after the beginning of sequences. For this reason, Table II-C reports the best results obtained during different runs of the same sequences, and omits results for which less than 30% of the trajectory was estimated.

Overall, there is no apparent winner among all the algorithms compared, in terms of lowest RMSE values, although for robustness, indicated by the proportions of sequences successfully processed, the approach of BASALT would be the most successful. Hence, this dataset can help the



Fig. 5. Dense maps produced by assembling LiDAR point clouds using pose estimates of OPEN_VINS, shown as a thin green line. (a-c) show maps produced on full sequences, the extent of the trajectories are visible in Table II. (d) shows a detail from the map of s3li_crater_inout, that highlights the inner rim of the Cisternazza crater and includes several examples of rock formations.

development of Visual SLAM approaches, considering the challenges provided by this type of environment, which are often missing in traditional datasets for Visual SLAM.

B. LiDAR Mapping

As Figs. 2 and 4 show, each LiDAR scan covers a small portion of the environment, which is located in front of the sensor setup within a relatively narrow Field-of-View. Furthermore, the geometry of the environment, which is characterized by soft slopes and small objects, does not provide structures of great extent in the vertical direction as it is, instead, very common in datasets targeted at autonomous driving [13], [44]. For these reasons, traditional LiDAR SLAM can not be employed successfully on this data. Compared to the stereo camera, however, where the depth uncertainty grows quadratically with the distance, the LiDAR sensor allows to measure accurately the shape of farther landscapes that might not be possible to reach, e.g., with a mobile system. It is, therefore, of great interest to develop SLAM pipelines where the LiDAR and visual sensing approaches are combined.

Meant purely as an example, we developed a simple mapping pipeline, that assembles LiDAR scans synchronized to pose estimated from a Visual SLAM approach. For this, we used OPEN_VINS. Fig. 5 shows four examples of LiDAR maps from the s3li_traverse_1, s3li_landmarks, s3li_traverse_2 and s3li_crater_inout sequences, respectively. Observing these maps, it is evident that the predominant geometries are slopes, with the exception of sparse and relatively small rocks and stones. Other more evident structures, i.e., big rock formations, are visible in Fig.5(b), which are re-observed twice. Fig 5(d), instead, shows the reconstruction of the rim of the Cisternazza crater, to an extent and detail that is not possible to achieve with a stereo camera alone. This dataset, therefore, intends to challenge visual-LiDAR SLAM algorithms to robustly perform pose estimation and mapping by developing novel approaches for place recognition that depart from traditional global LiDAR descriptors [45] or segmentation of structures with strong ground plane assumptions [46].

In order to illustrate what may lead to such methods, we employed the Instance Stereo Transformer (INSTR) [47], to detect instances of unknown objects in a scene (see Fig. 6 for



Fig. 6. Examples of instance segmentation of unknown objects, in this case small stones, using INSTR [47]. Shown are original left camera images (left column) and instance predictions as colored overlays (right column).

two examples). Note that the network has not been trained on this specific data. The segmentation masks can be used as an input for isolating objects in, otherwise complicated to process, LiDAR scans and accurately generate geometric models with the goal of realizing landmark-based [46] or metric-semantic [48] SLAM.

IV. CONCLUSIONS

We presented a dataset, recorded in a planetary analogue environment, that comprises stereo, LiDAR and inertial measurements, and allows to assess localization and mapping frameworks under challenging visual and structural appearance. We, furthermore, evaluated a variety of visual-inertial SLAM algorithms on the sequences, highlighting strengths and limitations, and provided examples of potential usecases for this data, that depart from traditional visual-LiDAR SLAM approaches. This dataset, therefore, provides a mean to evaluate and overcome limitations of traditional SLAM algorithms, that are usually developed and tested on datasets which lack elements of visual aliasing and ambiguous geometries, typical of extremely unstructured environments.

ACKNOWLEDGMENT

This work was supported by the Helmholtz Association, project ARCHES (contract number ZT-0033). Special thanks to Florian Auinger for his work on the system, to Lukas Meyer for the support on the mountain and to Emanuel Staudinger for the help on the D-GNSS setup.

REFERENCES

- Bresson, et al., "Simultaneous localization and mapping: A survey of current trends in autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 3, pp. 194–220, 2017.
- [2] Bürki, et al., "VIZARD: Reliable visual localization for autonomous vehicles in urban outdoor environments," in *IEEE Intelligent Vehicles* Symposium (IV), 2019, pp. 1124–1130.
- [3] Mascaro, et al., "Towards automating construction tasks: Large-scale object mapping, segmentation, and manipulation," *Journal of Field Robotics*, vol. 38, no. 5, pp. 684–699, 2021.
- [4] Shu, et al., "SLAM in the field: An evaluation of monocular mapping and localization on challenging dynamic agricultural environment," in *IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1761–1771.
- [5] Oliveira, et al., "Advances in agriculture robotics: A state-of-the-art review and challenges ahead," Robotics, vol. 10, no. 2, p. 52, 2021.
- [6] Sharafutdinov, et al., "Comparison of modern open-source visual SLAM approaches," arXiv preprint arXiv:2108.01654, 2021.
- [7] Dellenbach, et al., "What's in my LiDAR odometry toolbox?" arXiv preprint arXiv:2103.09708, 2021.
- [8] Shan and Englot, "LeGO-LOAM: Lightweight and ground-optimized lidar odometry and mapping on variable terrain," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4758–4765.
- [9] Giubilato, et al., "GPGM-SLAM: a robust SLAM system for unstructured planetary environments with gaussian process gradient maps," CoRR, vol. abs/2109.06596, 2021.
- [10] Le Gentil, et al., "Gaussian process gradient maps for loop-closure detection in unstructured planetary environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1895–1902.
- [11] Gelfand, et al., "Geometrically stable sampling for the ICP algorithm," in International Conference on 3-D Digital Imaging and Modeling (3DIM), 2003, pp. 260–267.
 [12] Geiger, et al., "Vision meets robotics: The KITTI dataset," The
- [12] Geiger, et al., "Vision meets robotics: The KITTI dataset," The International Journal of Robotics Research, vol. 32, no. 11, pp. 1231– 1237, 2013.
- [13] Maddern, et al., "1 year, 1000 km: The Oxford RobotCar dataset," The International Journal of Robotics Research, vol. 36, no. 1, pp. 3–15, 2017.
- [14] Choi, et al., "KAIST multi-spectral day/night data set for autonomous and assisted driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 3, pp. 934–948, 2018.
- [15] Wenzel, et al., "4Seasons: A cross-season dataset for multi-weather SLAM in autonomous driving," in DAGM German Conference on Pattern Recognition. Springer, 2020, pp. 404–417.
- [16] Sturm, et al., "A benchmark for the evaluation of RGB-D SLAM systems," in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012, pp. 573–580.
- and Systems (IROS), 2012, pp. 573–580.
 [17] Schubert, et al., "The TUM VI benchmark for evaluating visualinertial odometry," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1680–1687.
- [18] Schops, et al., "BAD SLAM: Bundle adjusted direct RGB-D SLAM," in IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
- [19] Handa, et al., "A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM," in *IEEE International Conference on Robotics* and Automation (ICRA), 2014, pp. 1524–1531.
- [20] Wang, et al., "TartanAir: A dataset to push the limits of visual SLAM," in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2020, pp. 4909–4916.
- [21] Furgale, et al., "The Devon Island rover navigation dataset," The International Journal of Robotics Research, vol. 31, no. 6, pp. 707– 713, 2012.
- [22] Leung, et al., "Chilean underground mine dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 16–23, 2017.

- [23] Miller, et al., "The visual-inertial canoe dataset," The International Journal of Robotics Research, vol. 37, no. 1, pp. 13–20, 2018.
- [24] Hewitt, et al., "The Katwijk beach planetary rover dataset," The International Journal of Robotics Research, vol. 37, no. 1, pp. 3–12, 2018.
- [25] Rogers, et al., "Test your SLAM! the SubT-Tunnel dataset and metric for mapping," in IEEE International Conference on Robotics and Automation (ICRA), 2020, pp. 955–961.
- [26] Bouman, et al., "Autonomous Spot: Long-range autonomous exploration of extreme environments with legged locomotion," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 2518–2525.
- [27] Rouček, et al., "DARPA subterranean challenge: Multi-robotic exploration of underground environments," in *International Conference on Modelling and Simulation for Autonomous Systems*, 2019, pp. 274–290.
- [28] Meyer, et al., "The MADMAX data set for visual-inertial rover navigation on Mars," Journal of Field Robotics, 2021.
- [29] Vayugundla, et al., "Datasets of long range navigation experiments in a Moon analogue environment on mount Etna," in *International* Symposium on Robotics (ISR), 2018, pp. 1–7.
- [30] Preston, et al., "CAFE concepts for activities in the field for exploration. TN2: The catalogue of planetary analogues," The Open University, Milton Keynes, UK, 2012.
- [31] Cramariuc, et al., "SemSegMap 3D segment-based semantic localization," in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021, pp. 1183–1190.
- [32] Schuster, et al., "Towards autonomous planetary exploration," Journal of Intelligent & Robotic Systems, vol. 93, no. 3, pp. 461–494, 2019.
- [33] Campos, et al., "ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multimap SLAM," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [34] Usenko, et al., "Visual-inertial mapping with non-linear factor recovery," IEEE Robotics and Automation Letters (RA-L) & International Conference on Intelligent Robotics and Automation (ICRA), vol. 5, no. 2, pp. 422–429, 2020.
- [35] Geneva, et al., "OpenVINS: A research platform for visual-inertial estimation," in IEEE International Conference on Robotics and Automation (ICRA), 2020.
- [36] Qin, et al., "A general optimization-based framework for local odometry estimation with multiple sensors," 2019.
- [37] Strobl, et al. DLR CalDe and DLR CalLab. http://www.robotic.dlr. de/callab. Institute of Robotics and Mechatronics, German Aerospace Center (DLR). Oberpfaffenhofen, Germany.
- [38] Furgale, et al., "Unified temporal and spatial calibration for multisensor systems," in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2013, pp. 1280–1286.
- [39] Zhou, et al., "Automatic extrinsic calibration of a camera and a 3D LiDAR using line and plane correspondences," in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2018, pp. 5562–5569.
- [40] Gong, et al., "3D LIDAR-camera extrinsic calibration using an arbitrary trihedron," Sensors, vol. 13, no. 2, pp. 1902–1918, 2013.
- [41] Tóth, et al., "Automatic LiDAR-camera calibration of extrinsic parameters using a spherical target," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 8580–8586.
- [42] Levinson and Thrun, "Automatic online calibration of cameras and lasers." in *Robotics: Science and Systems*, 2013.
- [43] Castorena, et al., "Autocalibration of LiDAR and optical cameras via edge alignment," in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2016, pp. 2862–2866.
- [44] Geiger, et al., "Are we ready for autonomous driving? the KITTI vision benchmark suite," in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012, pp. 3354–3361.
- [45] Kim and Kim, "Scan Context: Egocentric spatial descriptor for place recognition within 3D point cloud map," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4802– 4809.
- [46] Dubé, et al., "SegMap: 3D segment mapping using data-driven descriptors," arXiv preprint arXiv:1804.09557, 2018.
- [47] Durner, et al., "Unknown object segmentation from stereo images," in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021, pp. 4823–4830.
- [48] Tian, et al., "Kimera-multi: Robust, distributed, dense metric-semantic slam for multi-robot systems," *IEEE Transactions on Robotics*, 2022.