

AN ACCURATE REAL-TIME UAV MAPPING SOLUTION FOR THE GENERATION OF ORTHOMOSAICS AND SURFACE MODELS

Alexander Kern^{1*}, Philipp Fanta-Jende^{2*}, Philipp Glira², Felix Bruckmüller² and Christoph Sulzbachner²

¹ Technische Universität Braunschweig – Institute of Flight Guidance – a.kern@tu-braunschweig.de

² AIT – Austrian Institute of Technology – Center for Vision, Automation and Control – Assistive & Autonomous Systems
(phillipp.fanta-jende, philipp.glira, felix.bruckmueller, christoph.sulzbachner)@ait.ac.at

TC I/II – WG: UAS and Small Multi-sensor Platforms: Concepts and Applications (I/II)

KEY WORDS: UAV, SLAM, real-time, orthorectification, surface model, rapid mapping.

ABSTRACT:

UAVs have become an indispensable tool for a variety of mapping applications. Not only in the area of surveying, infrastructure planning and environmental monitoring tasks but also in time-critical applications, such as emergency and disaster response. Although UAVs enable rapid data acquisition per se, data processing usually relies on offline workflows. This contribution presents an accurate real-time data processing solution for UAV mapping applications as well as an extensive experimental and comparative study to the commercial offline solution Pix4D on the absolute accuracy of orthomosaics and digital surface models. We show that our procedure achieves an absolute horizontal and vertical accuracy of about 1 m without the use of ground control. The code will be made publicly available.

1. INTRODUCTION

In recent years, unmanned aerial vehicles (UAV) have become an important asset for rapid information retrieval for a wide array of applications, such as infrastructure inspection, environmental monitoring or disaster response (Ejaz et al., 2019; Erdelj et al., 2017; Kerle et al., 2019).

UAVs provide high agility, flexibility and fast data capture, whereas respective data processing chains, i.e. image orientation, 3D reconstruction and ortho-generation, are usually performed in post-processing. This traditional pipeline is suited for scenarios which require high-accuracy and excellent reconstruction quality. On the other hand, quick overviews of a scene, disaster or emergency response and other time-critical applications highly benefit from real-time capable processing workflows (Bu et al., 2016; Hein et al., 2019; Hinzmann et al., 2018).

This paper extends our previous work on this topic (Fanta-Jende et al., 2020; Kern et al., 2020) by integrating and comparing various state-of-the-art SLAM implementations (ORB-SLAM3 (Campos et al., 2020), OV²SLAM (Ferrera et al., 2021) and (Sumikura et al., 2019)) and their impact on the quality and accuracy of generated data products, i.e. orthomosaics and surface models.

2. RELATED WORK

Generating orthomosaics from a sequence of aerial images is a well-studied subject. Typically, a very distinct photogrammetric workflow is utilised that involves feature extraction and matching, camera pose estimation as well as sparse and dense reconstruction of the scene before the final orthomosaic is created. Open source and commercial implementations of such workflows are diverse (e.g. Pix4D, Agisoft Metashape, Colmap) and achieve high accuracies with the proper hardware and acquisition techniques involved. However, these

implementations are designed for post-processing applied to the images after a flight.

When time is crucial and maps are required to be obtained while the UAV is still in the air, the available solutions are limited. The work of Botterill et al. (2010) represents the traditional real-time approach avoiding costly computations of 3D camera motion aligning images into a mosaic with 2D projective homographies. This approach, however, assumes the ground to be planar, which is violated by surface elevation e.g. in low altitudes, mountainous areas or urban canyons resulting in distorted maps. To overcome this limitation Bu et al. (2016) used a modern visual SLAM to track the full camera movement and consequently project images into a common reference plane to create a global image mosaic. While this strategy is more robust when observing an elevated surface, the projection into a plane remains a simplification. In contrast, Hein et al. (2019) utilise an a-priori DEM to account for the elevation and rectify the images with a pinhole model. This reduces computational requirements significantly while still allowing for the generation of accurate maps. However, pre-existing DEMs are limited in availability and resolution (e.g. SRTM) while rapid changes due to e.g. natural disasters (earthquakes, floods, ...) or construction works can change the observed scene significantly. By tracking the camera using sensor fusion and reconstructing the surface during flight, Hinzmann et al. (2018) propose a solution that is able to provide the same 2D and 3D information as traditional offline photogrammetry while still achieving real-time performance. Our previous work OpenREALM (Kern et al., 2020) followed a similar approach but modularised important steps of the processing pipeline. Consequently, not only the camera trajectory, a digital surface model (DSM) and an orthomosaic are extracted in real-time but the framework can also be utilised as a testbed for modern visual SLAM and rapid 3D reconstruction applied to aerial mapping. Even though the initial results were promising, an extensive study on the accuracy of the final data products was still open and is now covered in this contribution.

* Corresponding authors

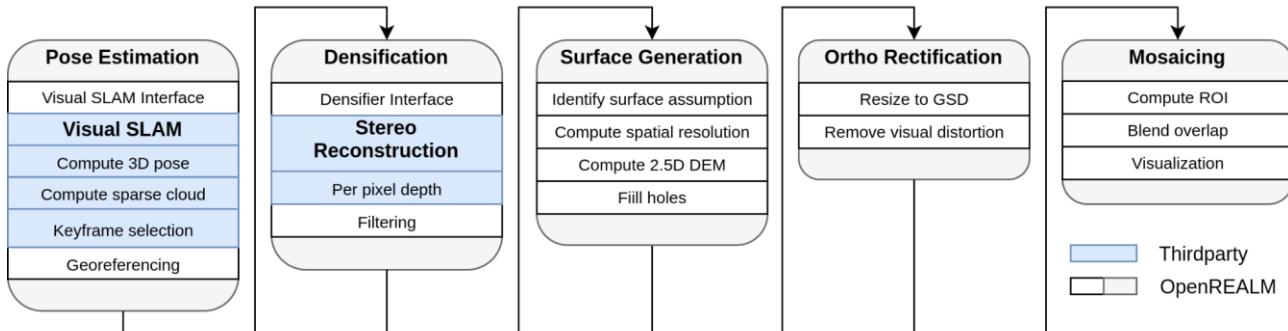


Figure 1 Structure of OpenREALM. It follows a pipeline design, separating different tasks into individual stages that are processed in parallel. An incoming frame moves from left (Pose Estimation) to right (Mosaicing).

3. METHODOLOGY

In the remainder of this work, we will analyse a variety of quality metrics for orthomosaics generated in real-time with OpenREALM. The workflow of the framework and its structure is shown in Figure 1 and explained briefly in the following. For further details please see Kern et al. (2020).

3.1 Processing modules

The input data consists of the GNSS observations, a heading estimate and the (aerial) images. This sensor data is fed into the system using a ROS¹ interface. Subsequently, the data is wrapped into a data frame which is passed into a series of individual stages. Each stage is encapsulated from the rest and processed in parallel allowing for fast overall performance of the pipeline.

3.2 Pose estimation

The first stage is the pose estimation, which tracks the camera motion by feeding the frames through an interface into a visual SLAM implementation of choice. Due to the monocular nature of the problem, scale is arbitrary. Consequently, the resulting pose is estimated in a local coordinate system. For every frame a local pose and an unsynchronised GNSS measurement are buffered.

Once a transformation between the local and geographic coordinate system can be computed using the Umeyama algorithm (Umeyama, 1991), this transformation is fixed and saved as a georeference. Depending on the user preference, this georeference can be refined with every new frame resulting in a higher absolute accuracy of the map, or it is fixed to prioritise the relative integrity of the map.

3.3 Densification

In the next stage, frames are passed to the densification module. At this point, the Plane Sweep Library (Häne et al., 2015) was integrated as an external framework to reconstruct depth from multiple views. As a result, dense depth maps are computed. The depth maps are filtered by checking the consistency of each depth value across multiple frames. Only if at least two other frames support the pixel depth hypothesis, it is considered valid.

3.4 Surface generation

Depending on the user settings and the computations of previous stages, frames that enter the surface generation are scanned for a surface assumption. If dense depth maps are available, a grid structure is created in which each cell contains an elevation value. The resolution of the grid depends on the average point to point distances of the dense cloud while corresponding elevation values are computed with the inverse distance weighting. This grid structure represents the digital elevation model that is further used for ortho-generation.

3.5 Orthorectification

Once an incremental DSM for the current input frame is generated, it is passed to the orthorectification module. Similar to the approach proposed by Hinzmann et al. (2018), a 3D point is created for every cell of the DSM and projected back into the camera using a pinhole model. Consequently, visual distortions induced by the elevated surface are removed and an incremental orthoimage is computed.

3.6 Mosaicing

In the final stage, incoming frames containing the DSM and orthoimage are fused. To this end, a region of interest that is affected by the latest frame update is extracted from the existing mosaic. Areas with no prior data are directly written. Those areas with existing data are blended by prioritising points with an incidence angle of 90° to achieve a higher degree of orthogonality. Textural blending, e.g. using seam cuts, is currently not implemented. Results are published and can be visualised in real-time as well.

3.7 Data products

All data products of traditional photogrammetry pipelines can be computed using OpenREALM. An overview with the processing stage and the corresponding data product is shown in Figure 2. The pose estimation module runs visual SLAM, consequently the 3D camera trajectory and a sparse point cloud can be exported. The densification module implements external stereo reconstruction in order to save depth maps for all incoming frames. The surface generation stage creates a 2.5D digital surface model, which is further used for orthorectification in the consecutive stage. Both, DSM and orthomosaic, are typically available in the GeoTIFF format. The mosaicing module wraps up all incremental data and is therefore designed to provide exports of all prior data products but at a global scale.

¹ ROS, <https://www.ros.org/>

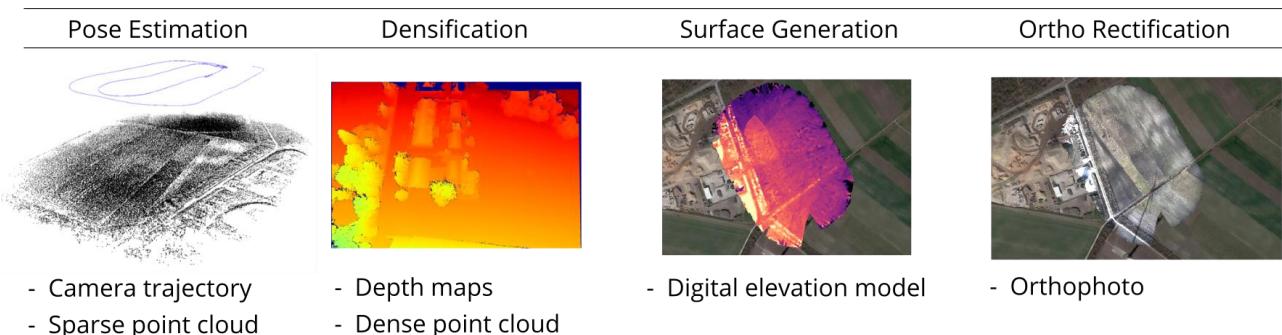


Figure 2 OpenREALM provides exports for most traditional photogrammetry data products. It allows for the extraction of the camera trajectory in 3D, sparse and dense point clouds, depth maps for all keyframes, 2.5D digital elevation models and orthomosaics, all of which are fully georeferenced.

3.8 Visual SLAM Varieties

OpenREALM is highly modular and can be easily extended. For this paper, we implemented interfaces for three state-of-the-art visual SLAM frameworks and analysed their feasibility for the aerial mapping scenario at hand.

The first is ORB-SLAM3 (Campos et al., 2020), whose predecessor set new standards in terms of accuracy and robustness in the computer vision community. It is a feature-based approach that reproduces the traditional photogrammetry workflow in high similarity. However, bundle adjustment is only applied to local subsets of keyframes to reduce computational load. Global integrity is achieved with a loop closing module that constantly checks for revisited areas using a bag-of-words model.

The second is OpenVSLAM (Sumikura et al., 2019), which is based on the same algorithms as ORB-SLAM but implements them in a more efficient way to further reduce computational load while tracking is improved overall.

At last, we integrated OV²SLAM (Ferrera et al., 2021) into the pipeline. It combines indirect, feature-based techniques for keyframes with a direct, photometric approach for frames in between. This enables robust tracking even when observing homogeneous surfaces while allowing for the usage of traditional optimisation and loop closing based on 2D features and a bag-of-words model.

4. EXPERIMENTAL DESIGN

4.1 Platform and hardware design

Our platform design aims at providing data acquisition in a quick and robust manner especially for disaster response scenarios using a fixed-wing platform. The data for this paper was acquired using an experimental UAV platform based on a Skywalker EVE-2000 equipped with a FLIR BFLY-U3-23S6C-C² featuring 2.3 MP. The camera is mounted on a gimbal compensating for deviations in the roll angle, which is especially useful for fixed-wing UAVs.

With a framerate of 8 images per second, high image overlap can be achieved at the standard velocity of 16 m/s of the platform. The system is equipped with a GNSS-RTK solution consisting of a Here3 receiver and a Here+ RTK base.

In its final stage, not only will the platform be changed to a more durable UAV, but the processing pipeline will be split up in an air and a ground segment. For instance, pose estimation could be performed on the UAV using the onboard processing capabilities of a NVIDIA Jetson TX2 while keyframes are sent to the ground station equipped with desktop-grade hardware for further and rather expensive processing stages, such as surface reconstruction.

4.2 Data acquisition

For the experimental setup presented in this paper, the data has been acquired in four separate flights at different altitudes: 80, 100, 120 and 150 metres above ground. Although these altitudes result in different ground sampling distances (GSD, see Table 1), all data products were processed with a grid size of 15 cm. The acquisition area is a rather flat and rural plot spanning roughly 400 by 300 metres south of Vienna, Austria. To ascertain the absolute accuracy of our mapping solution, 27 ground control points (GCP) have been surveyed using a GNSS rover.

Altitude [m]	Minimum GSD [cm]	Number of acquired images
80	5.43	1202
100	6.53	1113
120	7.96	947
150	10.05	911

Table 1 Specifics of the acquisition campaign; the minimum (theoretical) GSD has been computed based on the altitude and the camera specifications

² FLIR BFLY-U3-23S6C-C

<https://www.flir.eu/products/blackfly-usb3/?model=BFLY-U3-23S6C-C>

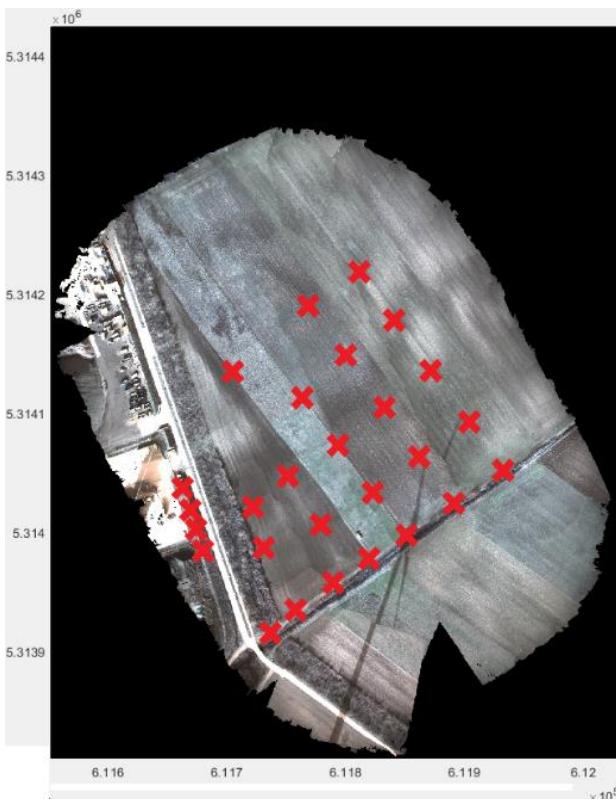


Figure 3 GCP locations in the acquisition area

4.3 Data preprocessing

To demonstrate the real-time capability of the presented processing pipeline, all data sets have been recorded in rosbags comprising the image frames, their extrinsic orientation parameters as well as precise timestamps. Since the entire pipeline runs on desktop hardware at the moment and will be split up in air and ground segment at a later stage in the project, rosbags are an ideal tool to simulate this real-time behaviour.

For generating reference data, the standard photogrammetric workflow in Pix4D was used. The settings were the following for all data sets: camera orientation: $\frac{1}{2}$ image size; point cloud generation: $\frac{1}{2}$ image size (standard), optimal point density and a minimum of observations per point set to three; DSM and ortho-generation: all settings set to standard except GSD of 15 cm. ROS timestamps were used to associate image frames and the respective exterior orientation parameters to export the images with EXIF headers. The interior parameters, which are required for the underlying SLAM pipeline of OpenREALM, were determined with Pix4D as well.

4.4 Data and methods for accuracy assessment

The aim of this study is to ascertain the absolute accuracy of the data products, i.e. orthomosaics in the horizontal dimension and digital surface models in the vertical dimension, generated by the proposed OpenREALM framework using different SLAM pipelines.

To this end, orthomosaics and DSMs were computed with three underlying SLAM pipelines (OpenVSLAM, ORB-SLAM3 and OV2SLAM) using data from different altitudes (80 m, 100 m, 120 m and 150 m). The SLAM pipeline has a tremendous impact on the feasibility of the DSM and orthomosaic

generation. For instance, lower flying altitudes lead to smaller frame overlaps which have an impact on the tracking capacity of the SLAM pipeline. Consequently, ORB-SLAM3 and OV2SLAM could not track visual features at 80 m and 100 m altitude, respectively, rendering a DSM and orthomosaic generation infeasible. Although ORB-SLAM3 offers the integration of inertial readings, it was used in vision-only mode to allow for a fair comparison.

Before evaluation, every visual SLAM was tuned inside the OpenREALM pipeline to maximise the overall performance and minimise the ground sampling distance while still maintaining a (near) real-time output. Resolution was prioritised over speed in order to be able to detect the GCPs in the final orthomosaic. All maps were fully processed within a minute after the respective mission was finished. Once a set of parameters was found for a SLAM, these were locked and used for all altitudes of the dataset in the evaluation. This ensures high quality results but prevents overfitting of the underlying models.

Pix4D was used to generate DSMs and orthomosaics for reference. Data from all altitudes has been processed without the use of GCPs in Pix4D. These data are closest for direct comparison to the real-time framework since resorting to ground control is unrealistic in time-critical mapping scenarios this procedure was designed for. To establish a baseline of the highest potential accuracy of acquired data, an orthomosaic and a DSM from the 120 m flight were computed in Pix4D using 10 GCPs.

Altitude	Real-time processing / Test data set		
	OpenVSLAM	ORB-SLAM3	OV2SLAM
80	Ortho/DSM	failed	failed
100	Ortho/DSM	failed	Ortho/DSM
120	Ortho/DSM	Ortho/DSM	Ortho/DSM
150	Ortho/DSM	Ortho/DSM	Ortho/DSM

Table 2 Generated data products in OpenREALM

Altitude	Offline processing / Reference data set	
	Pix4D without GCPs	Pix4D with 10 GCPs
80	Ortho/DSM	-
100	Ortho/DSM	-
120	Ortho/DSM	Ortho/DSM
150	Ortho/DSM	-

Table 3 Generated data products in Pix4D

The horizontal accuracy of the orthomosaics was assessed using all 27 GCPs as check points in the area. In the case of the Pix4D orthomosaic and DSM processed with ground control, the remaining 17 GCPs were used as check points. The vertical accuracy of the real-time data products was determined by sampling equally distributed points every 5 metres along the vertical and horizontal coordinate axis to establish a measuring grid. As a reference, a DSM from the 120 m data set computed by Pix4D with ground control is used.

5. RESULTS

This section presents quantitative accuracy results of the procedure as well as exemplary figures of the data products for visual comparison.

5.1 Real-time orthomosaic

The results are shown in the five tables below separated according to the underlying SLAM implementation and the reference data sets for comparison respectively. The reference data set processed without the use of ground control indicates an absolute accuracy of around 0.5 m with the exception of the data set acquired at an altitude of 100 m (Table 5). Using GCPs improves the absolute accuracy of the data set significantly as shown in Table 4.

Pix4D reference (with GCPs) - accuracy checked with 19 check points

Altitude	120	
RMSE	X	0.031
	Y	0.023
		-0.009
Mean	X	-0.002
	Y	-0.029
Sigma	X	0.022
	Y	-0.088
Min	X	0.037
	Y	0.036
Max	X	0.042
	Y	0.451

Table 4 Absolute accuracies of reference orthomosaic processed in Pix4D **with** ground control

Pix4D reference (without GCPs)				
Altitude	80	100	120	150
RMSE	X	0.654	1.211	0.411
	Y	0.395	1.33	0.335
Mean	X	0.515	-0.542	0.265
	Y	-0.208	0.774	-0.002
Sigma	X	0.41	1.102	0.319
	Y	0.342	1.098	0.341
Min	X	-0.247	-2.543	-0.344
	Y	-0.879	-1.288	-0.567
Max	X	1.268	1.826	0.955
	Y	0.357	2.729	0.608

Table 5 Absolute accuracies of the reference orthomosaic processed in Pix4D **without** ground control

It becomes evident that the absolute accuracy varies significantly depending on the SLAM pipeline. ORB-SLAM3 performs worst in our experiments (Table 6). Further studies will show whether this is the case if inertial readings are integrated in the procedure. The vision-only mode in conjunction with our pipeline reaches metre-grade accuracy at best. For the case of rapid mapping, however, these accuracies may be sufficient.

ORB-SLAM 3		
Altitude	120	150
RMSE	X	7.152
	Y	4.523
Mean	X	1.25
	Y	7.001
Sigma	X	3.272
	Y	4.369
Min	X	-0.581
	Y	1.459
Max	X	1.301
	Y	1.198
Min	X	1.128
	Y	5.013
Max	X	-2.647
	Y	1.229
Min	X	5.723
	Y	9.819
Max	X	2.022
	Y	7.911

Table 6 Absolute accuracies of the orthomosaic processed with ORB-SLAM3 as the underlying framework. Entries in bold are the lowest RMSE among the OpenREALM results in that dimension.

OpenVSLAM is the only SLAM framework which is able to track features reliably at 80 m altitude, although the absolute accuracy deteriorates in comparison to higher altitudes. OpenVSLAM reaches accuracy levels on par with the Pix4D data set at 100 m altitude (Table 7).

OpenVSLAM				
Altitude	80	100	120	150
RMSE	X	2.116	1.284	1.734
	Y	2.059	1.188	1.675
Mean	X	1.108	0.172	-0.703
	Y	1.296	0.267	0.403
Sigma	X	1.838	1.299	1.617
	Y	1.631	1.182	1.658
Min	X	-1.675	-1.691	-3.395
	Y	-1.517	-1.632	-2.2
Max	X	3.992	2.762	2.918
	Y	4.821	3.045	4.237

Table 7 Absolute accuracies of the orthomosaic processed with OpenVSLAM as the underlying framework. Entries in bold are the lowest RMSE among the OpenREALM results in that dimension.

The most recently published competitor OV²SLAM performs exceptionally well at 100 and 120 metres altitude reaching an absolute accuracy of around 1 m in both dimensions.

OV ² SLAM			
Altitude	100	120	150
RMSE	X 1.149	1.192	1.502
	Y 1.284	1.17	2.977
Mean	X -0.127	-0.684	0.921
	Y 0.449	0.349	2.551
Sigma	X 1.163	0.993	1.209
	Y 1.225	1.137	1.564
Min	X -2.746	-2.219	-0.728
	Y -1.944	-1.205	0.231
Max	X 1.707	0.851	3.281
	Y 2.473	2.868	6.17

Table 8 Absolute accuracies of the orthomosaic processed with OV²SLAM as the underlying framework. Entries in bold are the lowest RMSE among the OpenREALM results in that dimension.

Comparing the orthomosaics on a visual basis draws a different picture. Figure 4 depicts an image detail at the south-west edge of the scene comprising a straight road leading to a junction surrounded by greenery. The underlying SLAM framework does not only have an impact on the accuracy but also the consistency of the orthomosaic – the higher the relative accuracy, the better the representation of the scene.

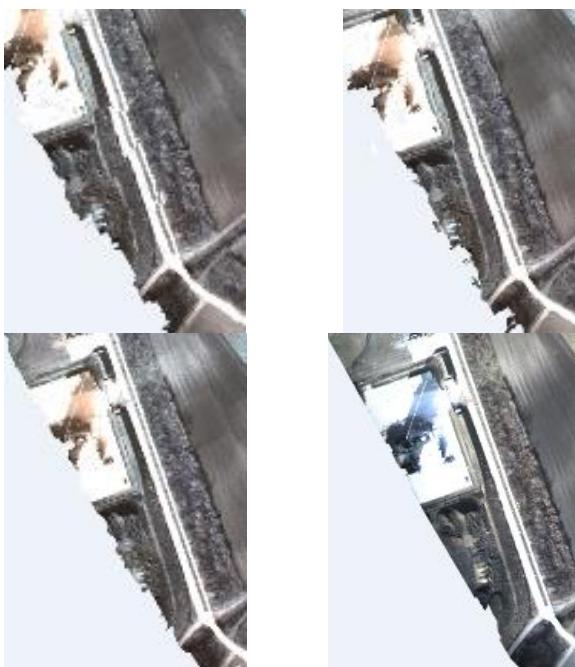


Figure 4 Comparison of an image detail between orthomosaics from the data set at 120 m. From the top left in clockwise order: OpenVSLAM, ORB-SLAM3, Pix4D reference data set processed with ground control and OV²SLAM

The processing pipeline based on OpenVSLAM (top left in Figure 4) produces a jittering appearance of the straight road while the ORB-SLAM3-based procedure (top right) does not result in a perfectly straight road either. OV²SLAM (bottom left), however, preserves the straight geometry of the road similar to the reference orthomosaic on bottom right. Moreover, the colouring differs between the real-time approaches and Pix4D. This is due to different blending procedures stitching the orthorectified image tiles together.

5.2 Real-time Digital Surface Model

A quantitative comparison of accuracy in height was performed based on equally sampled points with a distance of 5 metres across the entire scene (see Table 8). As a reference, the digital surface model generated by Pix4D using the 120 m data set with introduced ground control has been used.

The results show that our real-time approaches achieve similar vertical accuracies compared to the offline Pix4D processing pipeline without ground control. The procedure relying on ORB-SLAM3 fluctuates strongly as minimum and maximum errors are comparatively high. This indicates that the surface reconstruction is rather inconsistent.

On the other hand, OpenVSLAM reaches the lowest height error in average and the lowest standard deviation of the SLAM-based approaches. Again, OV²SLAM returns the most promising result among the three SLAM frameworks with an RMSE lower than the offline result.

The offline result has a very low standard deviation with a consistent mean and RMSE error. This may indicate that there is a general shift in height compared to the reference surface model. In the future, aerial laserscanning data could be used to perform more insightful understanding of height accuracy.

Comparison of height accuracy				
	Pix4D w/o GCPs	ORB- SLAM3	Open VSLAM	OV ² SLAM
RMSE Z	1.852	3.895	1.759	1.204
Mean Z error	-1.85	-1.028	0.006	-1.153
Sigma Z error	0.087	3.759	0.176	0.347
Minimum Z error	-2.029	-8.673	-4.915	-0.247
Maximum Z error	-1.616	12.287	4.698	-0.173

Table 9 Comparison of height accuracies between the reference surface model generated by Pix4D with ground control, the three real-time data products as well as the Pix4D reference data set without ground control

Figure 5 depicts the surface models generated by OpenREALM as well as the reference model by Pix4D. The figure shows that the surface reconstruction quality varies heavily. The OpenVSLAM-based procedure returns an uneven and rough surface which results in issues of preserving straight geometries (cf. Figure 4). The DSM based on the ORB-SLAM3 implementation tilts towards lower heights in the eastern part which indicates a drift in scale. The reconstruction using OV²SLAM as a basis achieves the best representation of the scene.

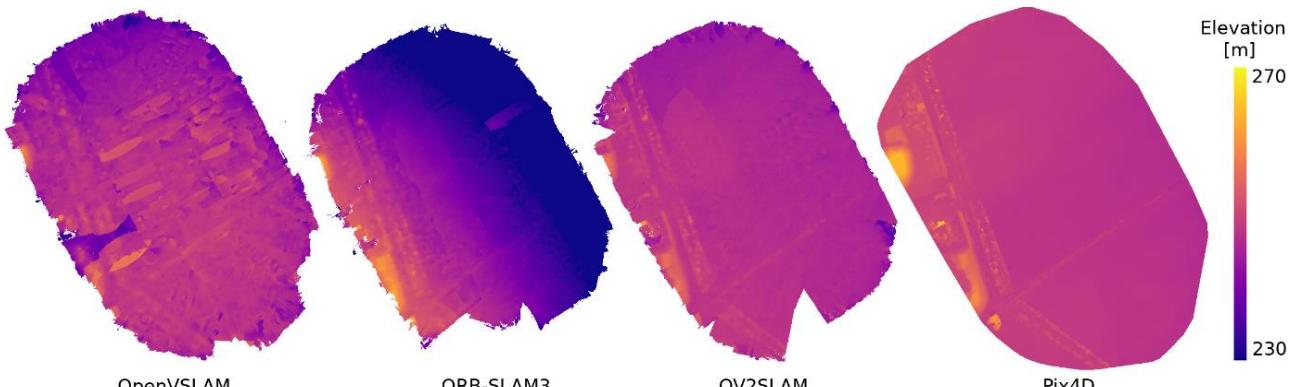


Figure 5 A visual comparison of digital surface models between the three SLAM-based real-time procedures and the reference data set processed with Pix4D

6. CONCLUSION

This paper presents a real-time mapping framework designed for UAVs. Real-time mapping is a pivotal tool for time-critical applications, such as disaster response and other emergency mapping tasks. We show that our framework enables generating overviews of a scene in a quick as well as accurate manner. Based on efficient SLAM implementations, an entire dense reconstruction pipeline and orthorectification procedure is able to perform high quality data processing in real-time.

In our experiments we compare three state-of-art SLAM implementations (ORB-SLAM3, OpenVSLAM and OV²SLAM) in conjunction with our real-time mapping pipeline OpenREALM. Our quantitative experiments indicate – depending on the underlying SLAM pipeline – that absolute accuracies of about 1 m horizontally and about 1 m vertically are well achievable, which is a promising result if compared to off-the-shelf offline reconstruction pipelines, such as Pix4D.

In our future work, we will integrate inertial readings into the pipeline to stabilise the trajectory estimation which in turn may potentially improve the robustness of the procedure at lower altitudes. Moreover, a thorough assessment of the height accuracy will be conducted with a LiDAR-based surface model. The code for our framework will be made publicly available.

REFERENCES

- Botterill, T., Mills, S., Green, R., 2010. Real-time aerial image mosaicing, in: International Conference Image and Vision Computing New Zealand. <https://doi.org/10.1109/IVCNZ.2010.6148850>
- Bu, S., Zhao, Y., Wan, G., Liu, Z., 2016. Map2DFusion: Real-time incremental UAV image mosaicing based on monocular SLAM, in: IEEE International Conference on Intelligent Robots and Systems. Institute of Electrical and Electronics Engineers Inc., pp. 4564–4571. <https://doi.org/10.1109/IROS.2016.7759672>
- Campos, C., Elvira, R., Rodríguez, J.J.G., Montiel, J.M.M., Tardós, J.D., 2020. ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial and Multi-Map SLAM. arXiv.
- Ejaz, W., Azam, M.A., Saadat, S., Iqbal, F., Hanan, A., 2019. Unmanned Aerial Vehicles enabled IoT Platform for Disaster Management. Energies 12, 2706. <https://doi.org/10.3390/en12142706>
- Erdelj, M., Natalizio, E., Chowdhury, K.R., Akyildiz, I.F., 2017. Help from the Sky: Leveraging UAVs for Disaster Management. IEEE Pervasive Comput. <https://doi.org/10.1109/MPRV.2017.11>
- Fanta-Jende, P., Steininger, D., Bruckmüller, F., Sulzbachner, C., 2020. A VERSATILE UAV NEAR REAL-TIME MAPPING SOLUTION FOR DISASTER RESPONSE – CONCEPT, IDEAS AND IMPLEMENTATION. ISPRS - Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci. XLIII-B1-2020, 429–435. <https://doi.org/10.5194/isprs-archives-XLIII-B1-2020-429-2020>
- Ferrera, M., Eudes, A., Moras, J., Sanfourche, M., Le Besnerais, G., 2021. OV2SLAM: A Fully Online and Versatile Visual SLAM for Real-Time Applications. IEEE Robot. Autom. Lett. 6, 1399–1406. <https://doi.org/10.1109/LRA.2021.3058069>
- Häne, C., Heng, L., Lee, G.H., Sizov, A., Pollefeys, M., 2015. Real-time direct dense matching on fisheye images using plane-sweeping stereo, in: Proceedings - 2014 International Conference on 3D Vision, 3DV 2014. Institute of Electrical and Electronics Engineers Inc., pp. 57–64. <https://doi.org/10.1109/3DV.2014.77>
- Hein, D., Kraft, T., Brauchle, J., Berger, R., 2019. Integrated UAV-Based Real-Time Mapping for Security Applications. ISPRS Int. J. Geo-Information 8, 219. <https://doi.org/10.3390/ijgi8050219>
- Hinzmann, T., Schönberger, J.L., Pollefeys, M., Siegwart, R., 2018. Mapping on the Fly: Real-Time 3D Dense Reconstruction, Digital Surface Map and Incremental Orthomosaic Generation for Unmanned Aerial Vehicles. Springer, Cham, pp. 383–396. https://doi.org/10.1007/978-3-319-67361-5_25
- Kerle, N., Gerke, Duarte, Vettrivel, 2019. UAV-Based Structural Damage Mapping: A Review. ISPRS Int. J. Geo-Information 9, 14. <https://doi.org/10.3390/ijgi9010014>
- Kern, A., Bobbe, M., Khedar, Y., Bestmann, U., 2020. OpenREALM: Real-time Mapping for Unmanned Aerial Vehicles.
- Sumikura, S., Shibuya, M., Sakurada, K., 2019. OpenVSLAM: A Versatile Visual SLAM Framework. MM 2019 - Proc. 27th ACM Int. Conf. Multimed. 2292–2295. <https://doi.org/10.1145/3343031.3350539>
- Umeyama, S., 1991. Least-Squares Estimation of Transformation Parameters Between Two Point Patterns. IEEE Trans. Pattern Anal. Mach. Intell. 13, 376–380. <https://doi.org/10.1109/34.88573>