

LiDAR-based Registration against Georeferenced Models for Globally Consistent Allocentric Maps

Jan Quenzel ^{a,b}, Linus T. Mallwitz ^a, Benedikt T. Arnold ^c, and Sven Behnke ^{a,b,d}

Abstract—Modern unmanned aerial vehicles (UAVs) are irreplaceable in search and rescue (SAR) missions to obtain a situational overview or provide closeups without endangering personnel. However, UAVs heavily rely on global navigation satellite system (GNSS) for localization which works well in open spaces, but the precision drastically degrades in the vicinity of buildings. These inaccuracies hinder aggregation of diverse data from multiple sources in a unified georeferenced frame for SAR operators.

In contrast, CityGML models provide approximate building shapes with accurate georeferenced poses. Besides, LiDAR works best in the vicinity of 3D structures. Hence, we refine coarse GNSS measurements by registering LiDAR maps against CityGML and digital elevation map (DEM) models as a prior for allocentric mapping. An intuitive plausibility score selects the best hypothesis based on occupancy using a 2D height map. Afterwards, we integrate the registration results in a continuous-time spline-based pose graph optimizer with LiDAR odometry and further sensing modalities to obtain globally consistent, georeferenced trajectories and maps.

We evaluate the viability of our approach on multiple flights captured at two distinct testing sites. Our method successfully reduced GNSS offset errors from up-to 16 m to below 0.5 m on multiple flights. Furthermore, we obtain globally consistent maps w.r.t. prior 3D geospatial models.

I. INTRODUCTION

Georeferenced maps are essential for modern SAR missions [1], [2]. From the initial planning stage to preserving a situational overview throughout the operation, first responders continuously compile status updates from diverse sources. Nowadays, UAVs are well-established for aerial overviews or imaging from difficult to reach viewpoints without endangering people’s life [3]–[10].

Automated UAV surveys require GNSS availability not only for integration with georeferenced maps, but also for safe flight in free airspace. Unfortunately, the precision drastically degrades in the vicinity of buildings due to reflections, shadowing and canyon effects [11]. This makes pure GNSS-based localization unreliable in urban areas where first responders need to fly in-between buildings. However, LiDAR-based odometry provides locally precise poses in the proximity of 3D structures, but is prone to accumulate drift over time and lacks georeferencing. In contrast, geospatial maps, e.g., CityGML and DEM, have accurate georeferenced poses for approximate building shapes and ground surfaces.

The registration of LiDAR maps with geospatial models promises to enhance localization in urban areas where GNSS

^a Autonomous Intelligent Systems Group, Computer Science Institute VI – Intelligent Systems and Robotics – and ^b Center for Robotics and Lamarr Institute for Machine Learning and Artificial Intelligence, University of Bonn, Germany; ^c Fraunhofer FIT, Germany; ^d Fraunhofer IAIS, Germany; quenzel@ais.uni-bonn.de

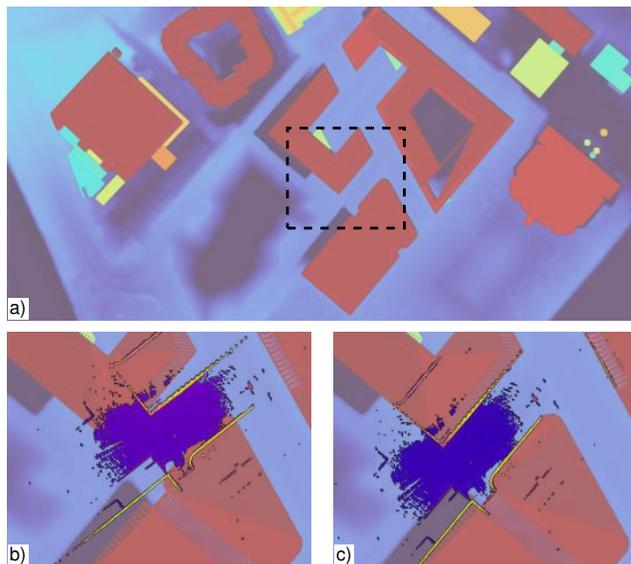


Fig. 1. Geospatial maps [a]) contain approximate building shapes and ground height. Inaccurate raw GNSS measurements impair the accuracy of georeferenced maps [b)]. We obtain a globally consistent map [c)] by registration against the geospatial model.

accuracy is low and LiDAR data is informative. Moreover, joint optimization, e.g., in a pose graph, reduces drift and provides strong priors for allocentric mapping.

To enable alignment of model and LiDAR data, CityGML data is combined with the corresponding DEM. The LiDAR odometry [12] processes scans to obtain local maps and poses. After semantic segmentation [13], we retain only walls and ground surfaces for the georeferencing. GNSS and IMU measurements initialize the approximate UAV pose — in general with an accuracy of few meters, resp. degrees. In a next step, we perform a grid search on the coarse horizontal offset and register the local map against the model at the offsetted poses using MARS [12]. A plausibility check determines the best matching result using an intuitive score from ray-traced occupancy with a 2D height map. The refined GNSS pose now georeferences the local LiDAR map and enables reliable registration of local maps for loop-closing.

For allocentric mapping, we directly optimize the B-spline knots of a continuous-time trajectory [14] using a pose graph. Odometry constraints connect scans to their local map. Additional constraints stem from relative transformation between local maps and preintegrated IMU [15]. We use the refined GNSS poses as anchors in our pose graph to obtain globally

consistent trajectories and allocentric maps with accurate georeferencing.

Our approach successfully reduces the GNSS error for the whole trajectory and even for local maps and single scans. In short, our contributions include:

- a refinement strategy for GNSS measurements using LiDAR registration against georeferenced 3D models,
- an intuitive plausibility score using height-based occupancy,
- an allocentric spline-based pose graph optimizer for continuous-time trajectories.

II. RELATED WORK

In recent years, several approaches have been developed to improve the quality of GNSS data using 3D models. Cappelle *et al.* [16] match RGB images against textured models and fuse GNSS, odometry and gyroscope in an Unscented Kalman-Filter (UKF). Wang *et al.* [17] tackle global localization on floor plans by matching corners against vertical edges within point clouds.

Zhang *et al.* [18] consider an autonomous driving scenario and combine LiDAR point clouds with GNSS. Instead of registering their LiDAR measurements against externally available 3D models, they make use of maps from previous runs that have been accurately georeferenced. The registration uses a deep neural network and measurements are directly fused within an Extended Kalman-Filter (EKF).

Lucks *et al.* [19] follow an approach similar to ours and register LiDAR scans against 3D models with the goal of mitigating the shortcomings of GNSS. As in our work, City Geography Markup Language (CityGML) is used in combination with a DEM. A major difference is their use of point-to-plane correspondences between raw scan and 3D model over long segments. Additionally, the transition between trajectory segments is interpolated to obtain a coherent map. Instead, we use adaptive surfel maps [12] and continuously register smaller segments if geometrical constraints allow successful registration.

Lv *et al.* [20] correct a continuous-time trajectory from a traditional pose graph of keyframes while maintaining the initial velocity prior to optimization. However, we directly optimize the trajectory and further include GNSS measurements as well as relative pose between previously refined local maps.

III. METHOD

Our method consists of several steps as shown in Fig. 2. We describe each step in the following and start with the description of the georeferenced model.

A. Georeferenced Model

Fortunately, many German state governments make geodetic data publicly available¹. Our georeferenced model combines a CityGML model with a DEM. The state of North

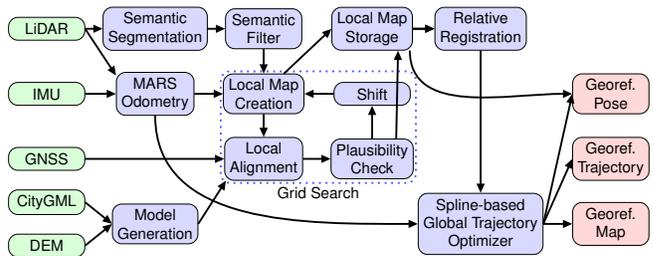


Fig. 2. System overview: Our refinement aligns small LiDAR maps against a geospatial model using GNSS for initialization. After pose graph optimization, our system outputs a globally consistent and georeferenced map and trajectory.

Rhine-Westphalia publishes both in 1 km² sized tiles². The CityGML contains the rough shape of buildings with Level of Detail (LoD)-2 whereas the DEM contains the grounds’ height in a grid with a 1.0 m² resolution. At first, we convert the CityGML model into a triangular mesh³ and subdivide the triangles [21] until their area is at most 0.1 m². Afterwards, we retain the triangle vertices. Similarly, we bi-linearly interpolate the DEM grid to 0.1 m² resolution and merge it with the sampled CityGML points for further processing.

If the CityGML model is not available or the model quality appears insufficient after visual inspection, we extract “roof” and cluster “contour” annotated points from Aerial Laser Scanning (ALS). After identifying “roof” points close to the contour, we extract the roof’s 2D α -shape without height using CGAL [22] and sample points in a line vertically from the roof down to the floor height.

For registration, MARS [12] derives a multi-resolution surfel map from the point cloud. Additionally, we compute a 2D height map where each cell stores the maximum height. This height map aids to assess the quality of the registration result with our plausibility score (Sec. III-D). Fig. 3 depicts the model of the Poppelsdorf Campus at the University of Bonn.

B. Scan Preparation

Our georeferenced model only contains ground and building surfaces. However, moving people, vegetation or other obstacles may be present in the actual LiDAR scans. Hence, we filter out clutter and retain only ground and building points using semantic segmentation [13], as shown in Fig. 4.

Filtered single scans \mathcal{P} are very sparse and more difficult to register. Thus, aggregating multiple scans into local maps \mathcal{W} creates more complete surfel maps. As measurements from the same position are redundant, the UAV needs to move more than τ m since the last accumulated scan.

C. Georeferenced Local Alignment

Given a local map and initial pose, we align its surfel map against the model’s surfel map with the registration

¹https://www.citygmlwiki.org/index.php?title=Open_Data_Initiatives_in_Germany

²<https://www.opengeodata.nrw.de/produkte/geobasis/>

³<https://github.com/citygml4j/citygml-tools>

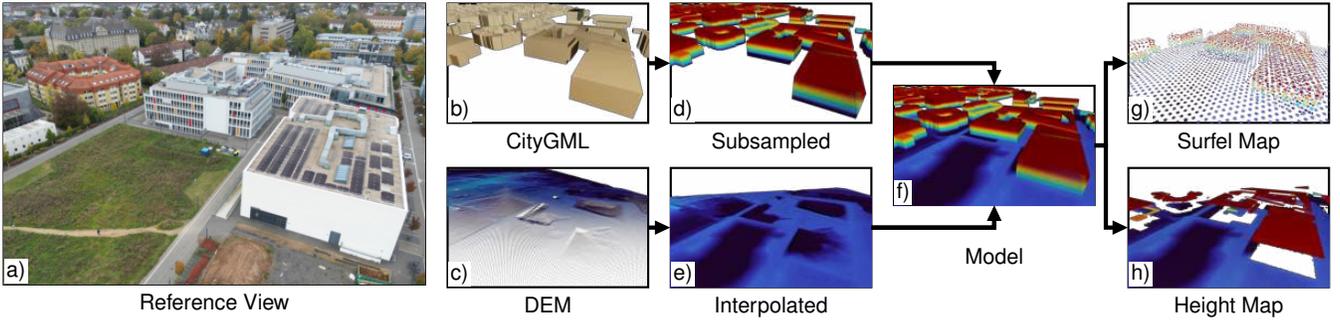


Fig. 3. Model Generation: A view [a]) on the Poppelsdorf campus at the University of Bonn for easier scene understanding. The CityGML data [b]) contains the rough polygonal building shape, while the DEM [c]) represents the ground surfaces. We combine the subsampled polygonal CityGML [d]) with the interpolated DEM [e]) for our model [f]). A surfel map [g]) is derived for registration [12] and a height map [h]) for our plausibility check.

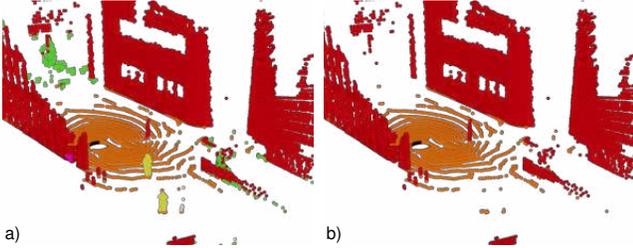


Fig. 4. A semantically annotated LiDAR scan with vegetation (green) and people (yellow) before [a]) and after filtering [b]).

of MARS [12] optimizing a single pose. Our initial guess for the local map pose stems from data provided by the UAV. A 3-axis magnetometer inside the IMU provides the orientation, whereas the GNSS supplies the approximate horizontal position and optionally the altitude. In proximity to the ground, we initialize the height from the model’s height map and the ultrasonic sensor of the UAV.

During tests, the horizontal GNSS position sometimes differed from the actual position by as much as 16 m. Such error exceeds the convergence radius for local registration and results in convergence to non-global local minima. We mitigate this with a grid search on the coarse horizontal shift and run the registration for each shifted initial pose. The grid covers the uncertainty range, e.g., with a resolution of 2 m, such that the local convergence basins overlap. Afterwards, a subsequent plausibility check (Sec. III-D) determines the best local minima.

In the absence of a magnetometer, the IMU accelerometer allows computation of roll and pitch whereas the grid search extends to combinations of horizontal shift and yaw.

D. Plausibility Check

Each alignment from the aforementioned grid search has to be checked for its plausibility to reject incorrect local minima. We propose a ray tracing-based plausibility score that operates solely on a discretized height map calculated from the model. The basic idea is to compare measured LiDAR rays with the corresponding projected rays in the model at the aligned pose.

After voxel filtering the scan points, we ray-trace horizontally using Bresenham’s line algorithm [23] from the sensor position \mathbf{o} towards the point \mathbf{p} in the height map. For every cell along the ray, we check that the ray $h_{\mathbf{p}_r}$ is above the height map $h_{\mathbf{p}_m}$. With the discretized ray distance $d_{\mathbf{p}}$ from \mathbf{o} to \mathbf{p} and the model distance $d_{\mathbf{p}_m}$ from \mathbf{o} to the first intersection ($h_{\mathbf{p}_m} > h_{\mathbf{p}_r}$), we compute a ray-score c_{ray} as follows:

$$c_{\text{ray}}(\mathbf{p}) = \min\left(\frac{d_{\mathbf{p}_m}}{d_{\mathbf{p}}}, 1\right) \in [0, 1]. \quad (1)$$

The score s_{ray} increases linearly with the measured ray distance up until the length of the projected ray in the model. Here, an upper limit of 1 ensures that measuring further, e.g., through windows, is neither penalized nor encouraged as the LoD-2 model only contains the rough shape of the facade and not its interior.

On its own, this would lead to incorrect results when comparing measurements on an open field with a building in the model. Hence, we introduce a binary hit score c_{hit} :

$$c_{\text{hit}}(\mathbf{p}) = \begin{cases} 1, & \text{if } (h_{\mathbf{p}_m} > (h_{\mathbf{p}} + \varepsilon)) \wedge (|d_{\mathbf{p}_m} - d_{\mathbf{p}}| < \vartheta), \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

Intuitively, it is implausible to measure an intersection ($h_{\mathbf{p}_m} > h_{\mathbf{p}}$) if there is no obstacle in the model map. At the same time, the endpoint should be close to the surface ($|d_{\mathbf{p}_m} - d_{\mathbf{p}}| < \vartheta$).

We obtain our plausibility score $s_{\mathcal{W}}$ for a local map as the mean over all scores with a linear combination of c_{ray} and c_{hit} with weight $w \in [0, 1]$:

$$c_{\mathbf{p}} = w c_{\text{ray}}(\mathbf{p}) + (1 - w) c_{\text{hit}}(\mathbf{p}), \quad (3)$$

$$s_{\mathcal{W}} = \frac{1}{|\mathcal{W}|} \sum_{\mathbf{p} \in \mathcal{W}} \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} c_{\mathbf{p}}. \quad (4)$$

As apparent from construction, the most plausible alignment from Sec. III-C should have $s_{\mathcal{W}}$ closest to one, which necessitates both criteria being close to one for all points.

Additionally, we compute the covariance’s condition numbers from model surfel map to local surfel map and vice versa to detect possible slippage of the translation components

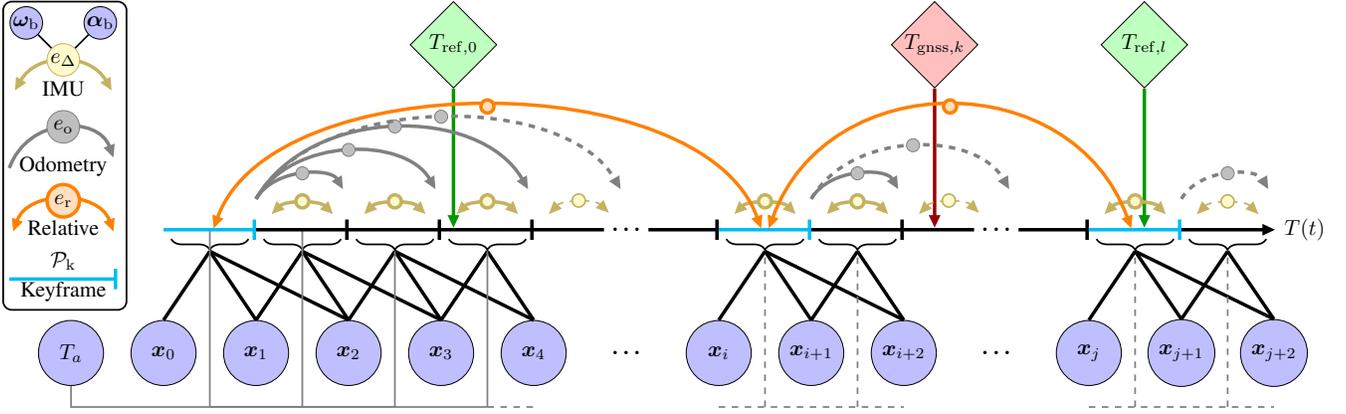


Fig. 5. Spline-based Pose Graph: We estimate a continuous-time B-spline trajectory [14] $T(t)$ with N knots $(\mathbf{x}_i, \dots, \mathbf{x}_{i+N-1} \in \mathcal{X})$ being active per scan. Raw and refined GNSS positions $(T_{\text{gnss}}, T_{\text{ref}})$ allow to georeference the UAV trajectory with an anchor pose T_a . Odometry constraints (e_o) connect each scan \mathcal{P} with the previous keyframe. Preintegrated IMU measurements (e_Δ) with optimizable biases (ω_b, α_b) enforce smoothness within a scan. Relative poses (e_r) between keyframes or scans enable loop closing.

during registration. Testing both directions ensures rejection even if associations differ. If κ is below a threshold τ and $s_{\mathcal{N}}$ is above γ , we accept the refined GNSS pose T_{ref} for the local map.

E. Spline-based Global Trajectory Optimization

In order to obtain a globally consistent map and trajectory, we build a pose graph $(\mathcal{V}, \mathcal{E})$ that represents the full UAV path using a continuous-time B-spline $T_{\mathcal{X}}(t)$ as in MARS [12]. In contrast to CLINS [20], we do not build a standard pose graph from some keyframes and optimize the keyframe poses only. Instead, our graph vertices \mathcal{V} directly contain the B-spline knots \mathcal{X} , the GNSS anchor pose T_a , IMU accelerometer and gyroscope biases (α_b, ω_b) , as shown in Fig. 5. The allocentric pose $T(t)$ at time t is given by:

$$T(t) = T_a T_{\mathcal{X}}(t). \quad (5)$$

We jointly minimize for all constraints $e \in \mathcal{E}$ the Mahalanobis error with distance \mathbf{d}_e and covariance Σ_e using a robust huber norm [24]:

$$\arg \min_{\mathcal{V}} \sum_{e \in \mathcal{E}} \rho_{\text{huber}}(\mathbf{d}_e^T \Sigma_e^{-1} \mathbf{d}_e). \quad (6)$$

A raw or refined GNSS pose T_{abs} with covariance Σ_{abs} provides an absolute constraint $e_a \in \mathcal{E}$ on $T(t)$:

$$\mathbf{d}_a = \text{Log}_{SE(3)}(T_{\text{abs}}^{-1} T_a T_{\mathcal{X}}(t)) \in \mathbb{R}^6, \quad (7)$$

using the logarithm map $\text{Log}_{SE(3)}$ [14]. Alternatively, this may be restricted to only the position \mathbf{p}_{abs} :

$$\mathbf{d}_{a,p} = (T_a \mathbf{p}_{\mathcal{X}}(t) - \mathbf{p}_{\text{abs}}) \in \mathbb{R}^3. \quad (8)$$

Odometry constraints $e_o \in \mathcal{E}$ connect from scan at time t_s towards the previous keyframe at t_k with pose T_o :

$$\mathbf{d}_o = \text{Log}_{SE(3)}(T_o^{-1} T_{\mathcal{X}}(t_k)^{-1} T_{\mathcal{X}}(t_s)) \in \mathbb{R}^6. \quad (9)$$

Preintegrated IMU measurements [15] $e_\Delta \in \mathcal{E}$ connect consecutive scans from $T_{\mathcal{X}}(t_{s-1})$ to $T_{\mathcal{X}}(t_s)$.

Additional relative pose constraints $e_r \in \mathcal{E}$ with $T_{\mathcal{X}}(t_1) \approx T_{\mathcal{X}}(t_0) T_{\text{rel}}$ stem from registration of time-wise or spatially neighboring local maps:

$$\mathbf{d}_r = \text{Log}_{SE(3)}(T_{\text{rel}}^{-1} T_{\mathcal{X}}(t_0)^{-1} T_{\mathcal{X}}(t_1)) \in \mathbb{R}^6. \quad (10)$$

Here, our refined GNSS poses T_{ref} aid in identifying spatially neighboring maps and initialize the relative pose T_{rel} for registration. We only add relative constraints if the initial translational distance for \mathbf{d}_r is smaller than 5% of the distance along the trajectory. This intuitively allows larger deviations for more distant loop-closures as errors accumulate over time.

Prior to optimization, we initialize yaw and horizontal 2D position of the GNSS anchor pose T_a by aligning [25] the refined GNSS positions with the corresponding spline positions. Empirically, we found this to provide a better initialization over longer segments than using a single pose since the IMU orientation might be slightly incorrect.

IV. EVALUATION

We recorded multiple UAV flights on different days at the Campus Poppelsdorf of the University of Bonn. Our UAV [26] is a modified DJI M210v2 equipped with an Intel NUC and an Ouster OS-0 128 LiDAR. The UAV has an external GNSS antenna by DJI to increase separation from the compute hardware and reduce possible interference.

Although, one can expect higher accuracy from RTK- and D-GNSS, our UAV is not equipped with either. Nonetheless, we recently experienced occasional offsets above 1 m and up to 4 m during static positioning tests with a Holybro UM982 RTK-GNSS due to canyoning.

We align the final maps against a georeferenced terrestrial LiDAR scanner (TLS) cloud of the campus using CloudCompare⁴. Afterwards, we compute the RMS positional error for the raw GNSS and refined GNSS measurements. Additionally, we evaluate the RMS positional error after

⁴<https://cloudcompare.org/>

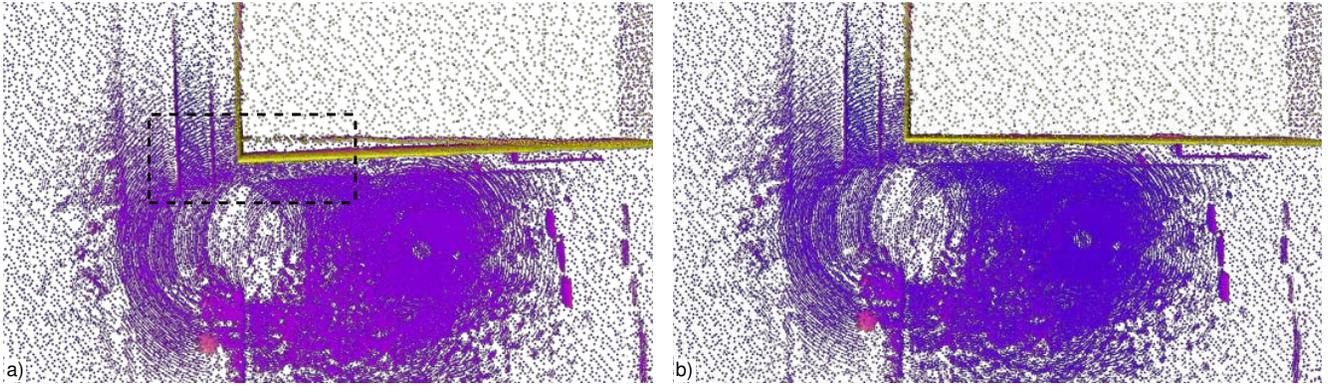


Fig. 6. Top-down closeup after optimization using only raw [a]) or refined GNSS [b]). Our refinement successfully corrected the GNSS offset and reduced the gap between height-colored model and local map.

TABLE I
RMSE EVALUATION W.R.T. GEOREFERENCED TLS

Sequence	raw GNSS [m]	ref. GNSS [m]	opt. w/ raw GNSS [m]	opt. w/ ref. GNSS [m]
11-13-36	1.896	0.039	1.808	0.033
11-17-39	0.813	0.575*	0.777	0.053
14-18-47	461.861 [†]	0.295	443.965	0.062
16-38-18	2.081	0.091	2.082	0.078
16-44-52	4.296	0.145	3.618	0.106
17-24-39	17.664	0.150	1.701	0.055

* due to outlier, w/o: 0.057 m. [†] contains mostly outlier.

optimization evaluated at the LiDAR scan time. The results are shown in Table I.

Fig. 6 highlights the difference between using raw and refined GNSS for georeferencing the reconstructed point cloud. The visible gap for raw GNSS vanishes after optimization when using refined GNSS. Fig. 7 further emphasizes this on sequence “16-38-18” for a larger section of the campus, showing clear improvement in aligning building walls.

The grid search for all offsets including alignment takes around 0.5 s using a radius of 8 m and 4 m step size. Hence, the search can run in parallel during normal operation since new local maps are only created sporadically. Our spline-based optimizer takes around 1.5 s for the 166 s long sequence “11-17-39” (Fig. 6).

We recorded 3 additional sequences at the abandoned Javelin Baracks in Elmpt, Germany during a forest fire training exercise by the fire brigade of the district Viersen. The ALS predominantly contains ground, plant, tree top and roof measurements with hardly any walls at a point density of ≈ 4 to 10 points per m^2 and an accuracy below 30 cm. In contrast to the previously used TLS, the ALS cloud exhibits too little overlap with our measurements to reliably constrain a reference alignment in lateral direction with sufficient accuracy. As a result, we showcase our results

in Fig. 8 and report the positional RMS distance for raw and refined GNSS w.r.t. our estimated correction. On average, our pipeline corrected the raw GNSS between 2.40 m and 3.06 m per sequence, whereas most refined GNSS positions are off by less than 10 cm.

V. CONCLUSION

We presented a novel approach to register local LiDAR maps against geospatial data to reduce GNSS offsets. A ray-tracing based score allows to select plausible refined GNSS poses. Our new spline-based global trajectory optimizer delivers globally consistent allocentric 3D maps. Our experiments showcased the effectiveness of our approach and successfully reduced the GNSS offset from multiple meters to below 0.5 m.

ACKNOWLEDGEMENT

We would like to thank Martin Blome and Lasse Klingbeil from the Institute of Geodesy and Geoinformation of the University of Bonn for the TLS reference map. Moreover, we would like to express our gratitude to the fire brigade of the district Viersen for the opportunity to participate in their forest fire training exercises. This work has been supported by the German Federal Ministry of Education and Research (BMBF) in the projects “Kompetenzzentrum: Etablierung des Deutschen Rettungsrobotik-Zentrums (E-DRZ)”, grant 13N16477, and “UMDenken: Supportive monitoring of turntable ladder operations for firefighting using IR images”, grant 13N16811.

REFERENCES

- [1] I. Kruijff-Korbayová, R. Grafe, N. Heidemann, A. Berrang, C. Hussung, C. Willms, P. Fetteke, M. Beul, J. Quenzel, D. Schleich, S. Behnke, J. Tiemann, J. Güldenring, M. Patchou, C. Arendt, C. Wietfeld, K. Daun, M. Schnaubelt, O. von Stryk, A. Lel, A. Miller, C. Röhrig, T. Straßmann, T. Barz, S. Soltau, F. Kremer, S. Rilling, R. Haseloff, S. Grobelny, A. Leinweber, G. Senkowski, M. Thurov, D. Slomma, and H. Surmann, “German Rescue Robotics Center (DRZ): A holistic approach for robotic systems assisting in emergency response,” in *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2021, pp. 138–145.

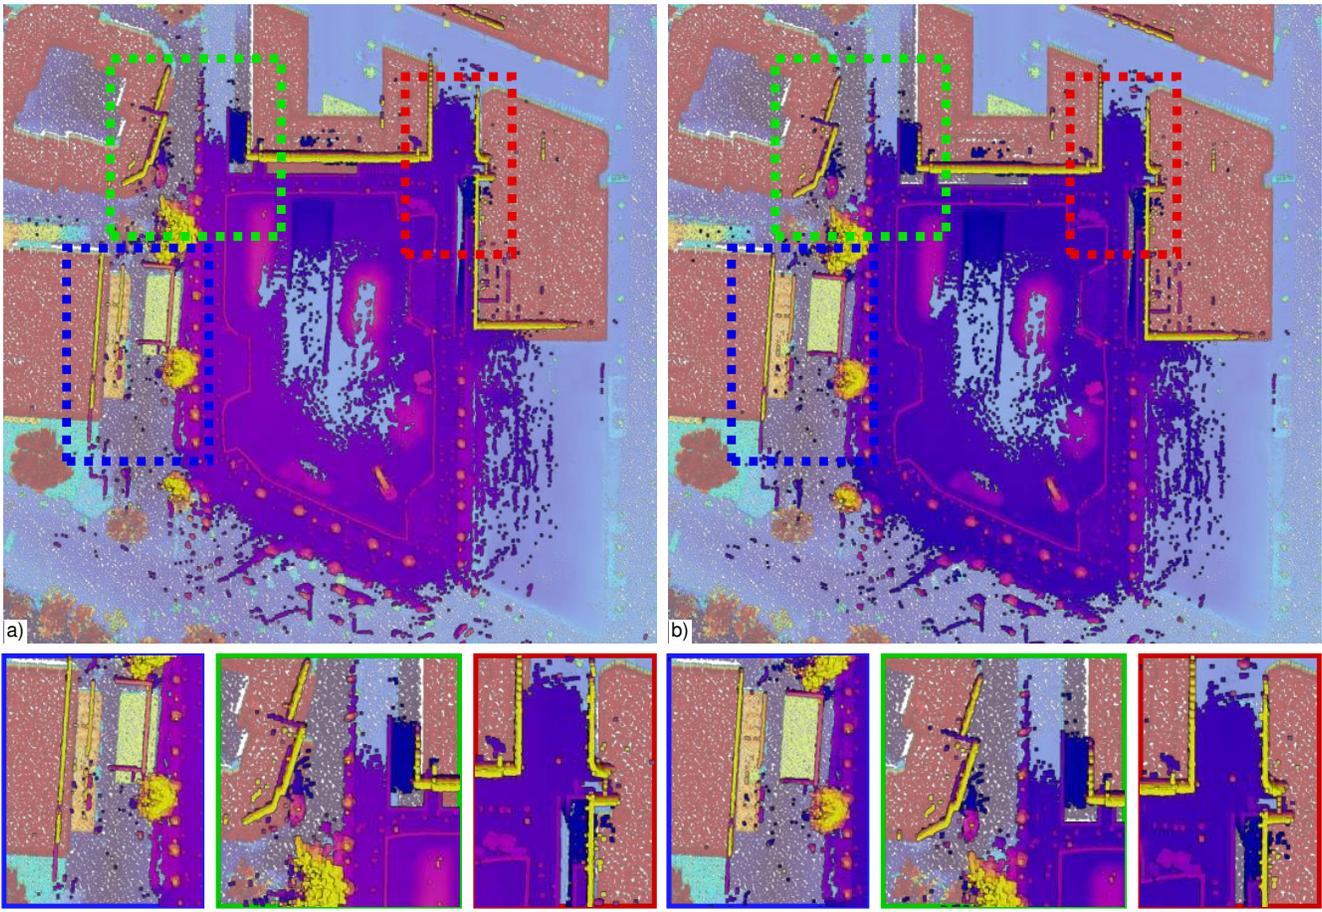


Fig. 7. Top-down view of the Poppelsdorf Campus at the University of Bonn after optimization using only raw [a)] or refined GNSS [b)]. For raw GNSS, facades (e.g., yellow lines) are mapped inside buildings (e.g., red areas). Our optimization successfully corrected the GNSS offsets and obtained a globally consistent map showing aligned facades w.r.t. the geospatial height-colored ALS map.

- [2] H. Surmann, K. Daun, M. Schnaubelt, O. von Stryk, M. Patchou, S. Böcker, C. Wietfeld, J. Quenzel, D. Schleich, S. Behnke, R. Grafe, N. Heidemann, D. Slomma, and I. Kruijff-Korbayová, “Lessons from robot-assisted disaster response deployments by the German Rescue Robotics Center task force,” *Journal of Field Robotics (JFR)*, vol. 41, no. 3, pp. 782–797, 2024.
- [3] S. MahmoudZadeh, A. Yazdani, Y. Kalantari, B. S. Ciftler, F. Aidarus, and M. O. Al-Kadri, “Holistic review of UAV-centric situational awareness: Applications, limitations, and algorithmic challenges,” *Robotics*, vol. 13, no. 8, p. 117, 2024.
- [4] M. Beul, M. Schwarz, J. Quenzel, M. Splietker, S. Bultmann, D. Schleich, A. Rochow, D. Pavlichenko, R. A. Rosu, P. Lowin, B. Scheider, M. Schreiber, F. Süberkrüb, and S. Behnke, “Target chase, wall building, and fire fighting: Autonomous UAVs of team nimbro at MBZIRC 2020,” *Field Robotics*, vol. 2, no. 1, pp. 807–842, 2022.
- [5] M. Lyu, Y. Zhao, C. Huang, and H. Huang, “Unmanned aerial vehicles for search and rescue: A survey,” *Remote Sensing*, vol. 15, no. 13, p. 3266, 2023.
- [6] R. A. Rosu, J. Quenzel, and S. Behnke, “Reconstruction of textured meshes for fire and heat source detection,” in *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2019, pp. 235–242.
- [7] I. Martinez-Alpiste, G. Golcarenenrenji, Q. Wang, and J. M. A. Calero, “Search and rescue operation using UAVs: A case study,” *Expert Systems with Applications*, vol. 178, p. 114937, 2021.
- [8] J. Quenzel, M. Nieuwenhuisen, D. Droeschel, M. Beul, S. Houben, and S. Behnke, “Autonomous MAV-based indoor chimney inspection with 3D laser localization and textured surface reconstruction,” *Journal of Intelligent and Robotic Systems (JINT)*, vol. 93, no. 1–2, pp. 317–335, 2019.
- [9] H. Hildmann and E. Kovacs, “Using unmanned aerial vehicles (UAVs) as mobile sensing platforms (MSPs) for disaster response, civil security and public safety,” *Drones*, vol. 3, no. 3, p. 59, 2019.
- [10] M. Beul, D. Droeschel, M. Nieuwenhuisen, J. Quenzel, S. Houben, and S. Behnke, “Fast autonomous flight in warehouses for inventory applications,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 4, pp. 3121–3128, 2018.
- [11] OpenStreetMap Wiki contributors, *Accuracy of GNSS data*. [Online; accessed 15-July-2024], 2022. [Online]. Available: https://wiki.openstreetmap.org/w/index.php?title=Accuracy_of_GNSS_data&oldid=2376050.
- [12] J. Quenzel and S. Behnke, “Real-time Multi-Adaptive-Resolution-Surfel 6D LiDAR Odometry using Continuous-time Trajectory Optimization,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [13] S. Bultmann, J. Quenzel, and S. Behnke, “Real-time multi-modal semantic fusion on unmanned aerial vehicles with label propagation for cross-domain adaptation,” *Robotics and Autonomous Systems*, vol. 159, p. 16, 2023.
- [14] C. Sommer, V. Usenko, D. Schubert, N. Demmel, and D. Cremers, “Efficient derivative computation for cumulative B-Splines on Lie groups,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [15] V. Usenko, N. Demmel, D. Schubert, J. Stückler, and D. Cremers, “Visual-inertial mapping with non-linear factor recovery,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 2, pp. 422–429, 2020.
- [16] C. Cappelle, M. E. El Najjar, F. Charpillat, and D. Pomorski, “Virtual 3d city model for navigation in urban areas,” *Journal of Intelligent and Robotic Systems (JINT)*, vol. 66, pp. 377–399, 2012.

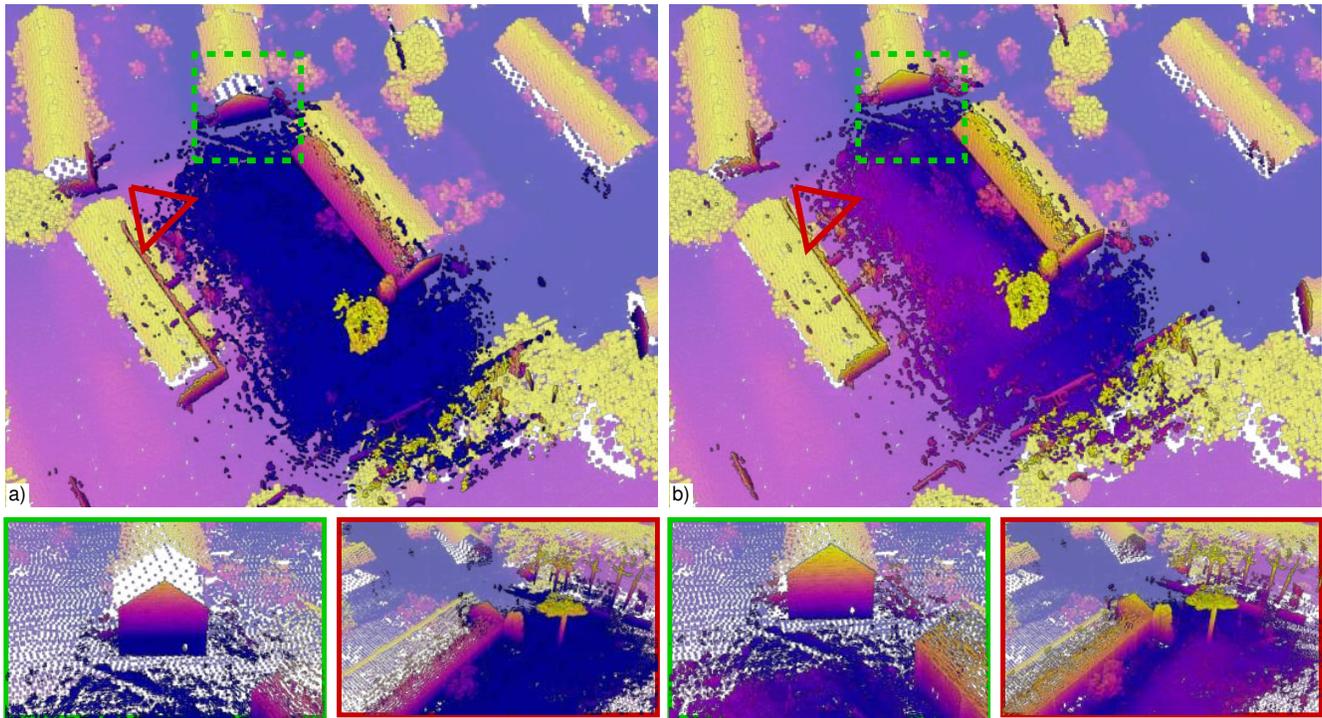


Fig. 8. Georeferenced maps after optimization using only raw [a]) or refined GNSS [b]) are overlaid on ALS reference points and color-coded by height. Sequence captured during a forest fire training exercise by the fire brigade of the district Viersen at the abandoned Javelin Baracks in Elmpt, Germany. ALS points are enlarged for better visibility. The gap between roof and measured wall (green rectangle) vanishes after optimization with refined GNSS. Large portions of the scene, as viewed from the red triangle, consist of vegetation which is non-represented in the CityGML and DEM.

- [17] X. Wang, R. J. Marcotte, and E. Olson, "GLFP: Global localization from a floor plan," in *IROS*, 2019, pp. 1627–1632.
- [18] J. Zhang, K. Khoshelham, and A. Khodabandeh, "Seamless vehicle positioning by lidar-GNSS integration: Standalone and multi-epoch scenarios," *Remote Sensing*, vol. 13, no. 22, p. 4525, 2021.
- [19] L. Lucks, L. Klingbeil, L. Plümer, and Y. Dehbi, "Improving trajectory estimation using 3D city models and kinematic point clouds," *Transactions in GIS*, vol. 25, no. 1, pp. 238–260, 2021.
- [20] J. Lv, K. Hu, J. Xu, Y. Liu, X. Ma, and X. Zuo, "CLINS: Continuous-Time Trajectory Estimation for LiDAR-Inertial System," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 6657–6663.
- [21] S. Seeger, K. Hormann, G. Häusler, and G. Greiner, "A sub-atomic subdivision approach," in *Vision Modeling and Visualization Conference (VMV)*, 2001, 77–86.
- [22] T. K. F. Da, "2D alpha shapes," in *CGAL User and Reference Manual*, 5.6.1, CGAL Editorial Board, 2024. [Online]. Available: <https://doc.cgal.org/5.6.1/Manual/packages.html#PkgAlphaShapes2>.
- [23] J. E. Bresenham, "Algorithm for computer control of a digital plotter," *IBM Systems journal*, vol. 4, no. 1, pp. 25–30, 1965.
- [24] P. J. Huber, "Robust Estimation of a Location Parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.
- [25] D. W. Eggert, A. Lorusso, and R. B. Fisher, "Estimating 3-d rigid body transformations: A comparison of four major algorithms," *Machine Vision and Applications*, vol. 9, pp. 272–290, 5 1997.
- [26] D. Schleich, M. Beul, J. Quenzel, and S. Behnke, "Autonomous flight in unknown GNSS-denied environments for disaster examination," in *International Conference on Unmanned Aircraft Systems (ICUAS)*, 2021.