Monocular Camera Localization in Prior LiDAR Maps with 2D-3D Line Correspondences

Huai Yu[1,2], Weikun Zhen[2], Wen Yang[1], Ji Zhang[2] and Sebastian Scherer[2]

Abstract— Light-weight camera localization in existing maps is essential for vision-based navigation. Currently, visual and visual-inertial odometry (VO&VIO) techniques are well-developed for state estimation but with inevitable accumulated drifts and pose jumps upon loop closure. To overcome these problems, we propose an efficient monocular camera localization method in prior LiDAR maps using direct 2D-3D line correspondences. To handle the appearance differences and modality gaps between LiDAR point clouds and images, geometric 3D lines are extracted offline from LiDAR maps while robust 2D lines are extracted online from video sequences. With the pose prediction from VIO, we can efficiently obtain coarse 2D-3D line correspondences. Then the camera poses and 2D-3D correspondences are iteratively optimized by minimizing the projection error of correspondences and rejecting outliers. Experimental results on the EurocMav dataset and our collected dataset demonstrate that the proposed method can efficiently estimate camera poses without accumulated drifts or pose jumps in structured environments.

I. INTRODUCTION

Accurate robot localization in urban environments is in great demand for autonomous vehicles. Since GPS localization is unstable without direct line-of-sight to the satellites, LiDAR-based localization modules are often used because of the accurate range measurements. Additionally, urban scenes have relative time-invariant geometric structures (e.g., buildings), thus one-time 3D map construction can be used for long-term localization. However, the expensive cost and heavy weight of LiDAR sensors limit its wide applications. Cameras and IMUs are low-cost, light-weight and commonly available sensors and current visual-inertial based pose estimation and mapping methods are well-developed for a variety of robot systems [1], [2], [3]. Nevertheless, state estimation methods that use only image features are prone to failures due to lighting or texture changes in the environment. Consequently, if camera localization modules can be associated with a prior 3D map, i.e., fuse the visual information with range measurements, there will be a great potential to use these light-weight and small camera modules for accurate localization in urban environments without a LiDAR sensor.

However, the fusion of image data with 3D point clouds is challenging due to appearance differences and modality gaps. Current approaches typically transfer the 3D data into 2D

[1]Huai Yu and Wen Yang are with the Electronic Information School, Wuhan University, Wuhan 430072, China {yuhuai, yangwen}@whu.edu.cn

[2] Huai Yu, Weikun Zhen, Ji Zhang and Sebastian Scherer are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA {weikunz, zhangji, basti}@andrew.cmu.edu
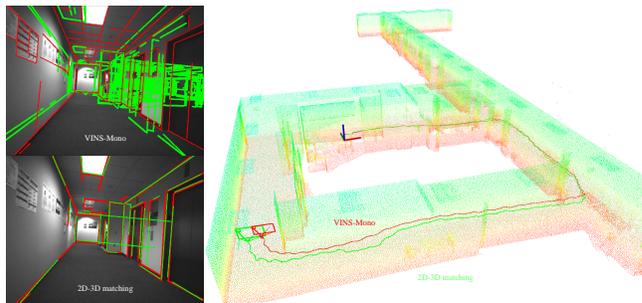
Fig. 1: The proposed monocular camera localization system in prior LiDAR maps of corridor. The right LiDAR map is colored by height. The red and green trajectories are the results of VINS-Mono [2] and ours, respectively. Top-left image shows the 3D line projections (green) using the estimated pose of VINS-Mono (with occlusions) and the extracted 2D lines (red), while the bottom-left image shows the 2D-3D correspondences using the pose estimation of the proposed method.

space or reconstruct 3D point clouds from 2D images to align data for pose estimation[4], [5]. Based on the characteristics of urban environments, our intuition lies on the fact that the major geometric structures, such as lines and planes, can be both captured in 3D maps and 2D images regardless of appearance differences and modality gaps. The direct 2D-3D geometric co-occurrence correspondence is more robust and precise than the association of domain-transferred data. Therefore, our purpose is to directly estimate the 2D-3D geometric line correspondences for the accurate and long-term camera localization.

In this work, we propose an approach for real-time light-weight monocular camera localization in prior 3D LiDAR maps using direct 2D-3D geometric line correspondences. We assume that a coarse pose initialization is given and focus on the pose tracking in maps, which follows the related works [6], [7]. For geometric concurrent feature extraction, 3D line segments are detected offline from LiDAR maps while robust 2D line segments are extracted online from video sequences. By employing the 6-DOF pose prediction from VIO, local visible 3D lines in field-of-view (FoV) are extracted and directly matched with 2D line features to obtain coarse 2D-3D line correspondences. Finally, the camera pose and 2D-3D matches are iteratively optimized by minimizing the projection error of correspondences and rejecting outliers.

The main contribution of this work is to estimate geo-

metric 2D-3D line correspondences for camera localization, which efficiently associates every keyframe with the prior LiDAR map. The geometric line correspondences are robust to appearance changes and suitable for camera localization in urban environments. Fig. 1 shows a camera image with 2D-3D line correspondences and estimated camera poses in the LiDAR map.

II. RELATED WORK

Vision-based localization in maps is to establish correspondences between 2D and 3D modalities for improving the localization robustness and accuracy [8]. To overcome the modality gaps and appearance differences, general approaches are using "intermediate products" to transfer the matching into the same space, i.e., in 2D space [9] or in 3D space [6], [10].

The first kind of camera localization in 3D maps is matching photometry in 2D image space. For most of the visual simultaneous localization and mapping (SLAM) methods [1], [2], sparse point cloud maps are reconstructed with photometry (or visual) descriptors, then loop closure detection is conducted by matching these descriptors to decrease drifts. In image-based localization methods [9], [4], large scale maps are reconstructed using SFM with visual features, these features are efficiently matched to yield 2D-3D correspondences for camera localization. Nevertheless, appearance and visual features are sensitive to illumination changes and light conditions, which make the correspondences unstable for long-term camera localization. Additionally, to handle the LiDAR maps without associated visual features, LiDAR appearance synthetic images are often used to directly match the camera images by normalized mutual information (NMI) [11] or normalized information distance (NID) [12]. Additionally, recent methods utilize colored point cloud projections [13] and synthesized depth images [14] to match with live images for camera localization. These localization methods are all trying to transfer the 2D-3D matching problem to be a 2D-2D matching problem.

The second strategy is matching geometry in 3D space to estimate camera poses. By using local bundle adjustment, [6] proposed to match a sparse reconstructed local 3D point clouds with given 3D LiDAR maps, which solves the scale estimation problem of monocular VO system and achieves online estimation of 6-DoF camera poses. Similarly in [15], 3D structural descriptors are used for matching LiDAR maps with sparse visual reconstructed point clouds. In [10], [7], dense local point clouds are reconstructed from a stereo camera to match to LiDAR maps, and then the matching results are loosely or tightly coupled into the VO and VIO system for optimizing camera poses. These localization methods by 3D registration obtain feasible results compared with vision-only based methods. However, the localization accuracy highly depends on local reconstruction performances. SFM suffers from scale problem and the reconstructed sparse points may not have correspondence in maps. Stereo reconstruction gives dense local point clouds

but mostly is time-consuming and does not scale well for long-range depth estimation.

Compared with the local point cloud matching methods, we aim at directly extracting 2D lines from a monocular camera to match with 3D lines from LiDAR maps for camera localization, which does not rely on SFM or stereo reconstruction modules. The global 2D-3D localization is known as a kidnapped robot problem because of the feature description gap and the non-convexity [16]. RANSAC-based and branch-and-bound strategies are often used to maximize 2D-3D inlier correspondences for global localization without a pose prior [16]. With the camera pose prediction from GPS or VO (VIO), we can obtain a pose prediction as the prior for local 2D-3D matching. SoftPosit [17] and BlindPnP [18] are conventional simultaneous 2D-3D correspondences and pose estimation approaches with the provided pose initialization. For the camera to LiDAR sensor calibration methods [19], [20], geometric features across images and point clouds are often matched to estimate the extrinsic transform between two sensors. Recently, some learning-based methods were proposed to directly register images to LiDAR-Map with GPS pose initialization [21], [22]. These methods suffer from either high computational complexity or unstable outputs, which still need explorations for realtime camera localization in 3D LiDAR maps.

III. PROPOSED METHOD

The proposed method simultaneously estimates 6-DoF camera poses and 2D-3D line correspondences in LiDAR maps. The correspondences are utilized to optimize camera poses by minimizing the 3D line projection error while refined camera poses can help reject outlier correspondences. As the preliminary to the online 2D-3D correspondence estimation, 3D line features are extracted offline on the large scale 3D LiDAR maps. At the same time, a coarse pose initialization is given for the first frame by the PnP solver [23] on manually labeled 2D-3D point correspondences. Then VINS-Mono [2] is utilized to predict the camera motion between adjacent keyframes. With the predicted poses, local 3D lines in camera field-of-view (FoV) are extracted and directly matched with the online extracted 2D lines from image sequences. Finally, camera poses and 2D-3D correspondences are iteratively updated. The pipeline is shown in Fig. 2.
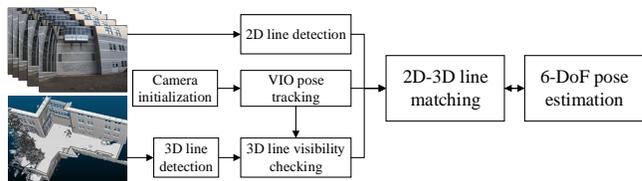

Fig. 2: Pipeline of the proposed camera localization method

A. 2D and 3D line extraction

In urban environments, the geometric structures are often represented by line segments and planes. We use a

segmentation-based 3D line detection method [24] to extract 3D lines from LiDAR maps. The general idea is to cluster point clouds into plane regions and use contour line fitting to obtain 3D line segments. This method is efficient and robust for large scale unorganized point clouds. Although it needs time to process millions of points, the 3D lines of all maps are extracted only once before we start tracking.

For 2D line extraction, we want to extract the major geometric 2D lines which are consistent with 3D lines and robust to noises. It is challenging in urban scenes because the substantial texture noises yield fragmented 2D line segments and some geometric edges are invisible in 2D images on the color-homogeneous structures (e.g., white wall). Many state-of-the-art line segment detection (LSD) methods have been presented in computer vision [25], [26], where traditional hand-craft methods are with high efficiency for online running on CPU. However, the detected lines are fragmented and noisy, an example is shown in Fig. 3(a). These kinds of fragmented and noisy features can produce a lot of 2D-3D matching outliers. Considering the line completeness and robustness to noisy, we finally employ a learning-based LSD [26] which uses the attraction field map (AFM) to transfer the LSD problem to be a region coloring problem. For each pixel p in images, the model first learns a 2D vector $\mathbf{a}(p)$ from the pixel to its nearest point p' on the nearest line segment.

$$\mathbf{a}(p) = \boldsymbol{p}' - \boldsymbol{p} \qquad (1)$$

where $\boldsymbol{p} = (x, y)$ is the coordinate of pixel p. Then an AFM $\{\mathbf{a}(p), p \in I\}$ is generated by encoding the pixels to line associations. During the training step, ground-truth 2D lines are transferred to AFM representation and then the network directly learns a model to minimize the similarity between the output and the true AFM. For testing, AFM representations are first generated from images, then a squeeze module [25] is used to iteratively group $\boldsymbol{v}(p) = \boldsymbol{p} + \mathbf{a}(p)$ belonging to the same line to fit line segments. Fig. 3(b) shows the detected line segments of the example image. It shows good consistency with geometric 3D structures and robustness to texture noises.
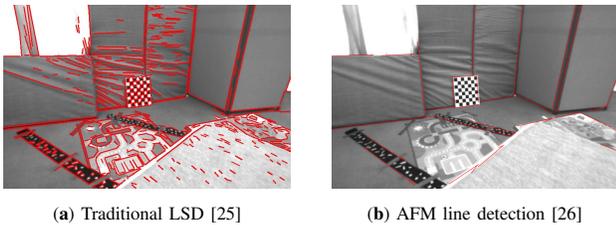

<table>
<tr><td>(a) Traditional LSD [25]</td><td>(b) AFM line detection [26]</td></tr>
</table>

Fig. 3: Comparison of different methods for 2D line segment detection.

B. 2D-3D line matching

For a single frame, the main steps to obtain 2D-3D correspondences consist of initial camera pose prediction,

visible 3D lines collection, and individual 2D-3D line correspondence estimation. Here the extraction of 3D lines in FoV helps improve the efficiency because local 3D lines in FoV are very limited compared with all 3D lines in the 3D map. Considering the occlusion checking is difficult to conduct on only 3D lines map, we keep all the 3D lines in FoV without discarding occluded lines.

For an image at time t, the corresponding pose estimation from VINS-Mono is denoted as $\bar{\mathbf{P}}_t \in SE(3)$, and the updated pose using 2D-3D correspondences is denoted as $\mathbf{P}_t \in SE(3)$. By using the estimated pose from 2D-3D correspondences of the last frame $\mathbf{P}_{t-1}$ and the camera motion from VINS-Mono $\mathbf{T}$, the updated pose $\hat{\mathbf{P}}_t$ can be computed

$$\begin{aligned} \mathbf{T} &= \bar{\mathbf{P}}_t \cdot (\bar{\mathbf{P}}_{t-1})^{-1}, \\ \hat{\mathbf{P}}_t &= \mathbf{T} \cdot \mathbf{P}_{t-1}. \end{aligned} \qquad (2)$$

With the pose prediction $\hat{\mathbf{P}}_t$, the local 3D lines $\{L_t\}$ in FoV can be extracted based on the two endpoint projections for efficiency. For checking the visibility of a 3D point $\boldsymbol{P}$ in FoV,

$$\boldsymbol{p}_t = \begin{bmatrix} x & y & w \end{bmatrix}^T = \mathbf{K} \cdot \hat{\mathbf{P}}_t \cdot \boldsymbol{P} \qquad (3)$$

where $\mathbf{K}$ is the camera intrinsic matrix. If $w > 0$, $0 \leq \lfloor x/w \rfloor < column$ and $0 \leq \lfloor y/w \rfloor < row$, the 3D point is in FoV. For a 3D line segment, the visibility checking is more complicated. Based on the visibility of two endpoints, there are three cases for validation:
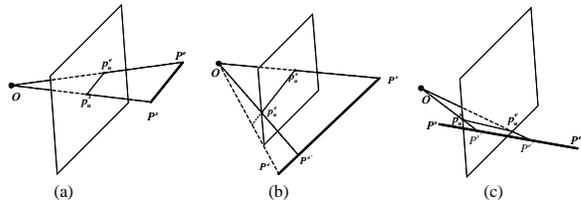

Fig. 4: 3D line visibility checking

- If both two endpoints are in FoV (Fig. 4(a)), we keep the whole 3D line segment as a local visible feature.
- When only one endpoint is in FoV (in Fig. 4(b)), we iteratively sample new 3D points on the 3D line from the visible point by 0.1 ratio of the line length and check the visibility of the new sample points. The generated subset 3D line segment with the longest length in FoV is stored as a local visible feature.
- When both two endpoints are out of FoV but a subset is in FoV, as shown in Fig. 4(c), we can also sample points to extract a subset. However, most of the invisible map lines are in this case, we discard all the 3D segments with two endpoints out of FoV for efficiency.

With the predicted camera pose $\hat{\mathbf{P}}_t$, 2D lines $\{l_t = (\boldsymbol{p}_t^s, \boldsymbol{p}_t^e)\}$ are directly matched with the local 3D lines in FoV, where $\boldsymbol{p}^s$ and $\boldsymbol{p}^e$ are the start and end points of a 2D line, respectively. For each possible 2D-3D correspondence, we use a 3D vector to measure the similarity, 2D angle distance θ, the distance d of two 3D endpoint projections

to the corresponding infinite 2D lines, and the overlapped length $len_{overlap}$ of the finite 2D lines with the 3D line projection.

$$dist(L_i, l_j) = \{\theta, d, len_{overlap}\} \tag{4}$$

The projection of a 3D line $L = (\boldsymbol{P}^s, \boldsymbol{P}^e)$ on the image plane is $l_u = (\boldsymbol{p}_u^s, \boldsymbol{p}_u^e)$ (Eq. 3), which is matched with the detected 2D line l_t. The normalized orientation of 2D lines is denoted as $\boldsymbol{v} = (\boldsymbol{p}^e - \boldsymbol{p}^s)/||\boldsymbol{p}^e - \boldsymbol{p}^s||$. Then the 2D angle distance θ can be computed by

$$\theta = \arccos(\boldsymbol{v}_t^T \cdot \boldsymbol{v}_u). \tag{5}$$

Assume the parametric representation of a extracted 2D line l_t is $Ax + By + C = 0$. The distance d can be computed by

$$d = \frac{|Au_x^s + Bu_y^s + C| + |Au_x^e + Bu_y^e + C|}{\sqrt{A^2 + B^2}}. \tag{6}$$

By using the point-to-line projection points, the overlapped length $len_{overlap}$ between 3D line projections l_u and the detected 2D lines l_t is

$$\begin{aligned} \alpha_* &= \arg\min_{\alpha \in [0,1]} ||\boldsymbol{p}_t^s + \alpha \cdot (\boldsymbol{p}_t^e - \boldsymbol{p}_t^s) - \boldsymbol{p}_u||_2^2, \\ len_{overlap} &= |\alpha_*^s - \alpha_*^e| \cdot ||\boldsymbol{p}_t^e - \boldsymbol{p}_t^s||_2^2, \end{aligned} \tag{7}$$

where α_*^s and α_*^e are corresponding to the projection points of $\boldsymbol{p}_u^s$ and $\boldsymbol{p}_u^e$ respectively. Here α^* provides an unified representation of point to finite line distance. Then for each extracted 2D line l_t^i, brute-force searching strategy is used to find a 3D line whose distance $\theta < \theta_0$ and $d < d_0$. Thus we obtain a set of coarse 2D-3D line correspondences for further pose estimation.

C. Pose optimization and correspondence refinement

For a single frame, the camera pose can be optimized by minimizing the point-to-infinite-line distance of the two 3D endpoint projections to corresponding 2D line distance. The Lie algebra of the estimated camera poses $\mathbf{P}_t$ are denoted as $\boldsymbol{\xi}_t$. The coefficient vector of the infinite 2D line can be $\mathbf{H} = \begin{bmatrix} A & B & C \end{bmatrix}$. The object function is to minimize the projection errors between all the 2D-3D correspondences:

$$\begin{aligned} \boldsymbol{\xi}_t^\star &= \arg\min_{\boldsymbol{\xi}_t} \sum_{i=1}^M d_i \\ &= \arg\min_{\boldsymbol{\xi}_t} \frac{1}{2} \sum_{i=1}^M \frac{||\mathbf{H}_i \cdot \mathbf{K} \exp(\boldsymbol{\xi}_t) L_i||_2^2}{\sqrt{A_i^2 + B_i^2}}, \end{aligned} \tag{8}$$

where L_i contains two endpoints, M denotes the number of correspondences. It is formulated as a non-linear least-squares problem. We initialize the camera pose with the predicted pose $\hat{\mathbf{P}}_t$ from VINS-Mono. With Lie algebra, the typical L-M algorithm can find the optimal camera pose $\mathbf{P}_t$.

However, single frame 2D-3D correspondence observations are not robust enough for online camera localization. When the 3D lines in FoV are limited or parallel to each other in 3D space, the 2D-3D correspondences cannot constrain the 6-DoF pose. Additionally, even the
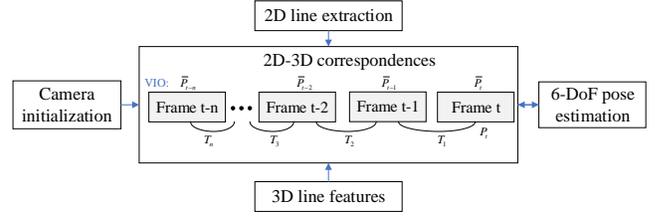

Fig. 5: Pose estimation in the sliding window

correspondences are enough for pose estimation, the geometric localization noises of both 2D and 3D lines will make the estimation jitter around the true pose. To solve these problems, a sliding window is utilized to add more previous correspondence observations for optimizing the current pose $\mathbf{P}_t$ (as shown in Fig. 5). Assuming the camera motion estimated from VINS-Mono for two adjacent keyframes is accurate, which is reasonable because the drifts of VINS-Mono on two adjacent keyframes are small. $\mathbf{T}_t^1, \mathbf{T}_t^2, \cdots, \mathbf{T}_t^n$ are denoted as camera motions between two adjacent keyframes (4×4 matrix including R and t). The camera pose $\mathbf{P}_{t-n}$ of the previous n-th keyframe can be derived by the current estimation $\mathbf{P}_t$ and camera motion

$$\mathbf{P}_{t-n} = (\mathbf{T}_t^n)^{-1} \cdots (\mathbf{T}_t^2)^{-1}(\mathbf{T}_t^1)^{-1}\mathbf{P}_t. \tag{9}$$

Then all previous 2D-3D correspondences in the sliding window can be "visible" in current frame. In Eq. 9, the previous camera pose $\mathbf{P}_{t-n}$ is related to the current pose $\mathbf{P}_t$, while $(\mathbf{T}_t^1\mathbf{T}_t^2 \cdots \mathbf{T}_t^n)^{-1}$ is a constant pose transformation, its Lie algebra is denoted as $\Delta\boldsymbol{\xi}_t^n$. Thus the pose optimization function is

$$\boldsymbol{\xi}_t^\star = \arg\min_{\boldsymbol{\xi}_t} \frac{1}{2} \sum_{n=0}^N \sum_{i=1}^M \frac{||\mathbf{H}_i^n \cdot \mathbf{K} \exp(\boldsymbol{\xi}_t + \Delta\boldsymbol{\xi}_t^n) L_i^n||_2^2}{\sqrt{A_i^{n\,2} + B_i^{n\,2}}}, \tag{10}$$

where N is the number of previous frames in the sliding window. Therefore, more observations make the estimation more robust to outliers, and the motion constraints smooth the pose trajectory. When using the sliding window, an unbalance problem will arise due to the different numbers of 2D-3D correspondences for each frame. To equalize the contribution of each frame correspondences and improve the efficiency, a threshold for the maximum number of correspondences is set to discard the 2D-3D correspondences with short overlap distance $len_{overlap}$. We employ Ceres Solver [27] to implement the optimization. After an optimized camera pose is obtained, the 2D-3D correspondences can be re-estimated by the updated pose to reject outliers. It follows the 2D-3D line matching in Section. III-B with more restrict thresholds (e.g., $\theta_0 = 0.8 * \theta, d_0 = 0.8 * d_0$). Then a more accurate camera pose can be updated with the new correspondences. After several iterations, both camera pose and 2D-3D correspondences can be optimized.

IV. EXPERIMENTAL RESULTS

The proposed method was tested on two different real-world dataset. The first experiment was conducted on the public available EuRoC MAV Dataset [28] with ground-truth trajectories. Then, we performed real experiments on our dataset collected by a Realsense D435i camera under varying conditions to validate the performance.

A. Implementation

Considering the 3D LiDAR maps are often large scale with a big data volume, we first down-sample the point clouds using CloudCompare [29] with a constant space resolution (1-5cm in experiments). Then the 3D line extraction based on [24] is conducted to obtain 3D lines. Since the arbitrary pose initialization in maps is a kidnapped robot problem for global 2D-3D correspondence estimation, a coarse pose initialization is given for the first frame by the PnP solver [23] on manually labeled 2D-3D point correspondences (at least 4 pairs, we use 6 pairs for robustness). For the learning-based 2D line extraction, we directly use the trained U-Net model [26] from Wireframe dataset [30] and modify it for online line segment detection. The testing platform is a desktop with Intel Core i7-4790K CPU, 32GB RAM, and an Nvidia GeForce GTX 980Ti GPU. The GPU is only used for 2D line detection.

During the 2D-3D correspondence estimation, the angle distance threshold θ_0 is set as 10 degrees to constrain the 3D projections to be almost parallel with corresponding 2D lines. Then the point to line distance threshold d_0 is set around $20 \sim 30$ pixels to collect matching pairs. If the number of correspondences for an image is less than a threshold (set as 8 empirically), it is identified unstable therefore the camera pose is predicted by camera motion only. If the number exceeds a threshold, we discard the correspondences with short $len_{overlap}$ distance for efficiency (empirically $M_0 = 40$, window size $N = 10$, there will be at most 440 correspondences in the sliding window). In the experiments, 2 or 3 iterations of optimization are often enough to obtain stable camera pose and 2D-3D correspondences.

Since our method is based on monocular visual-inertial odometry without loop closure, we compare it with the two versions of VINS-Mono [2], i.e., with or without loop closure. Loop closure helps for reducing the overall absolute trajectory error and mapping, but the refinement of the past poses in the loop does not help for realtime localization. Additionally, it causes pose shifts for current pose which have side effects for robot navigation.

B. Result on EuRoC MAV Dataset

The EuRoC MAV Dataset [28] is a visual-inertial datasets collected onboard a UAV. The datasets contain stereo images (20Hz), synchronized IMU data (200Hz), accurate ground-truth trajectories about 70 meters and a scan LiDAR map. The subset room data consist of 2 LiDAR maps wherein 3 video sequences each. 2D-3D correspondence results and estimated trajectory in LiDAR maps are shown in Fig. 6. All the 3D lines in FoV are projected to the image plane

without occlusion checking. From the top left image, we can observe that the projections of 3D lines (in green) are shifted by using the estimated pose of VINS-Mono(odom). Then by iteratively updating 2D-3D correspondences and camera poses, stable correspondences are obtained with a more accurate pose. The position drifts are greatly reduced and stable 2D-3D correspondences can be estimated.
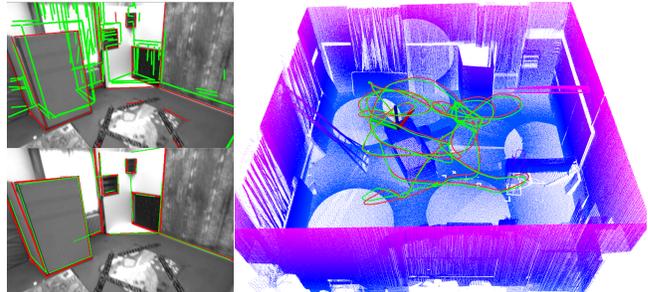

Fig. 6: Camera localization results on the EuRoC MAV dataset. The top left image shows the extracted 2D lines (red) and the projected 3D lines (green) using VINS-Mono(odom) (with occlusions), bottom left shows the 2D-3D correspondences using our method. The right image shows the estimated trajectory (green) aligned with ground-truth (red) in LiDAR maps.

TABLE I: ATE RMSE [31] results over 5 runs on the EuRoC MAV Dataset.

Dataset	VINS-Mono (odom)[2]	VINS-Mono (loop)[2]	2D-3D matching
V1_01_easy	0.147	**0.075**	0.089
V1_02_medium	0.153	0.127	**0.069**
V1_03_difficult	0.301	**0.152**	0.173
V2_01_easy	0.275	**0.099**	0.166
V2_02_medium	0.221	0.135	**0.132**
V2_03_difficult	0.726	**0.324**	0.635

For quantitative analysis, the beginning 200 estimated poses of each sequence are used for trajectory alignment with ground-truth [32]. The absolute trajectory error (ATE) [31] results are shown in Table I, it is clear that the 2D-3D correspondences improve the pose estimation accuracy compared with only odometry. The V2 room has more noises making the 2D-3D correspondences sometimes unstable for pose optimization. This is the reason why the improvements for the V2 room are less significant than the results of the V1 room. The worst case is that no stable 2D-3D correspondence is available and the final estimations follow the odometry. Furthermore, our method shows competitive results compared with VINS-Mono(loop). While the loop closure optimizes the past poses in the loop and produces a pose jump for the current pose. For the realtime localization purpose, the refinements of the poses in the past make no sense and the pose jumps have side effects for navigation. Our method always estimates the current poses in the sliding window, which greatly reduces the drifts and does not have the issue of pose jumps.

The average relative pose errors (RPE) [32] are shown in Table II and Fig. 7, which are used to show the growth of position error with the trajectory length. The accumulated drift of VINS-mono(odom) is growing up with travel length. While the error of our method keeps small and stable along the way, which is related to the accuracy of 2D and 3D line localization.

TABLE II: RPE RMSE [31] over different segment lengths

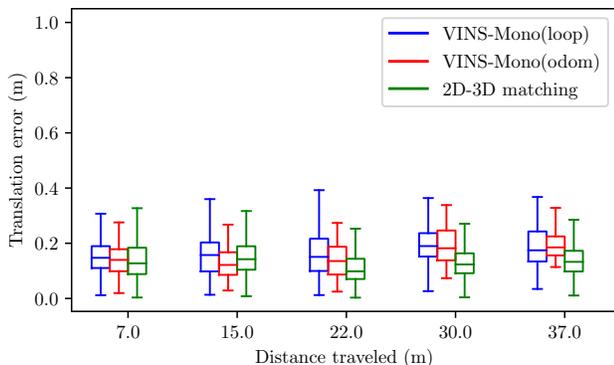| Segment Length(m)	VINS-Mono (odom)[2]	VINS-Mono (loop)[2]	2D-3D matching
7	**0.150**	0.166	0.173
15	0.161	**0.154**	0.172
22	0.184	0.174	**0.139**
30	0.200	0.1840	**0.158**
37	0.212	0.190	**0.168**

Fig. 7: Boxplot of the relative trajectory errors

C. Evaluation on our dataset

To further evaluate our method under various environments, we tested it on our own collected data of indoor corridors and outdoor buildings. An Intel RealSense D435i camera is used to collect synchronized images and IMU data. The left global-shutter imager captures monocular image sequences (640×480 pixels images at 30Hz, the IR projector turned off) with synchronized IMU data (200Hz). The LiDAR maps are obtained by registering several scans of a FARO scanner focus3D S, as shown in Fig. 8. Both the indoor corridors (Fig. 8a) and Smith Hall (Fig. 8c) have a lot of occlusions, while the NSH building (Fig. 8b) is much simpler with fewer occlusions. For these experiments, the trajectories are in the same pattern to run a complete round and return to the start points to see the position drifts. The loop closure is not stably detected for all runs, so the results of VINS-Mono with loop closure are not discussed.

The indoor corridor's results are shown in Fig. 1. The results of two outdoor buildings are shown in Fig. 9. Considering that we do not have ground-truth trajectories, the estimation accuracy validation is shown in the following two ways. For the qualitative analysis, the 3D line features are projected to overlap with 2D lines using the estimated poses

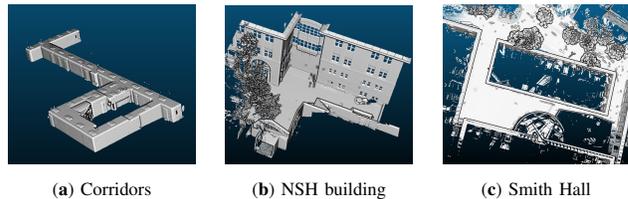(a) Corridors   (b) NSH building   (c) Smith Hall
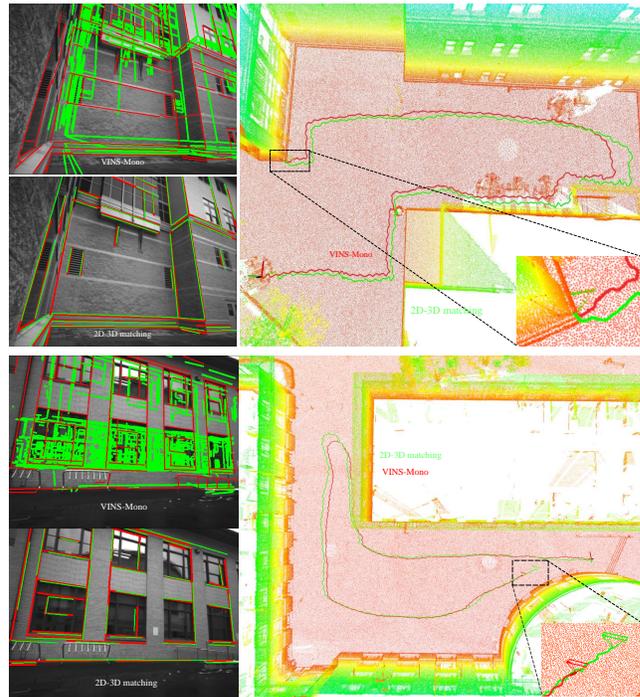
Fig. 8: Three urban scene LiDAR maps.

Fig. 9: Camera localization results in outdoor environments

of VINS-Mono(odom) and our method. For VINS-Mono odometry, it can be observed that the 3D line projections (green lines) are shifted and scaled by inaccurate camera poses in the top left images. While using our method, the pose estimations provide more accurate and stable 2D-3D structure correspondences in the bottom left images. Additionally, we can observe obvious misalignment between the two trajectories. More online correspondence and localization results are shown in the video [1].

For the accuracy evaluation, we pick 5 frames along the trajectories and use the PnP solver to estimate the ground-truth camera poses on manually labelled 10 pairs of 2D-3D point correspondences. The 5 frames are sampled along the trajectories on different times ($t = \beta * t_0$, t_0 is the total running time) and the RMSE positions along 5 runs are used as the final ground-truths. The position errors are shown in Table III. The accumulated drifts increase a lot along the trajectories for VINS-Mono(odom). While our method greatly improves the localization accuracy with the assistance of stable 2D-3D line correspondences. The localization errors

[1]https://youtu.be/H80Bnxm8IPE

keep small along the whole trajectories. Another interesting observation is that the accumulated error can drift back if we keep the direction of the system and backward to the start position, which is shown in the results of VINS-Mono on NSH building.

TABLE III: Localization errors in man-made environments

β	Corridors		NSH building		Smith Hall	
	VINS-	2D-3D	VINS-	2D-3D	VINS-	2D-3D
$(t = \beta t_0)$	Mono	matching	Mono	matching	Mono	matching
0.1	0.125	**0.109**	**0.137**	0.152	0.354	**0.180**
0.3	0.596	**0.128**	0.319	**0.127**	0.507	**0.154**
0.5	0.438	**0.094**	1.203	**0.170**	1.057	**0.201**
0.8	0.575	**0.120**	0.352	**0.156**	1.245	**0.156**
1.0	0.705	**0.112**	0.621	**0.132**	2.176	**0.178**
Length(m)	130		95		120	

In terms of efficiency, VINS-Mono does not use map information and can be customized set for output frequency. However, with different setting frequencies, the odometry results change a lot. We select the most stable one at 15 Hz. Then for the estimation of 2D-3D correspondences and the camera poses, it costs about 0.01 seconds on average for each keyframe. Since the 3D line extraction is offline before the system starts, 2D line detection can run at 25Hz on 640×480 images, our method can run at about $13 \sim 15$ Hz for all the scenarios.

V. CONCLUSION

In this paper, we presented a novel monocular camera localization approach in prior LiDAR maps of structured environments. With the 3D geometric lines from LiDAR maps and robust online 2D line detection, our method efficiently obtains coarse 2D-3D line correspondences based on the camera motion prediction from VINS-Mono. The pose optimization with 2D-3D correspondences greatly reduces the pose estimation drifts of VIO system without using visual-revisiting loop closure. Both qualitative and quantitative results on real-world datasets demonstrate that our method can efficiently obtain reliable 2D-3D correspondences and accurate camera pose in LiDAR maps. As future work we intend to enhance the robustness of 2D-3D correspondences on inaccurate pose predictions, such as directly using the pose of the last frame as the prediction.

ACKNOWLEDGMENT

REFERENCES

[1] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.

[2] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.

[3] K. Sun, K. Mohta, B. Pfrommer, M. Watterson, S. Liu, Y. Mulgaonkar, C. J. Taylor, and V. Kumar, "Robust stereo visual inertial odometry for fast autonomous flight," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 965–972, 2018.

[4] Y. Feng, L. Fan, and Y. Wu, "Fast localization in large-scale environments using supervised indexing of binary features," *IEEE Transactions on Image Processing*, vol. 25, no. 1, pp. 343–358, 2016.

[5] J. Engel, T. Schöps, and D. Cremers, "Lsd-slam: Large-scale direct monocular slam," in *European Conference on Computer Vision*, 2014, pp. 834–849.

[6] T. Caselitz, B. Steder, M. Ruhnke, and W. Burgard, "Monocular camera localization in 3d lidar maps," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 1926–1931.

[7] X. Zuo, P. Geneva, Y. Yang, W. Ye, Y. Liu, and G. Huang, "Visual-inertial localization with prior LiDAR map constraints," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3394–3401, 2019.

[8] N. Piasco, D. Sidibé, C. Demonceaux, and V. Gouet-Brunet, "A survey on visual-based localization: On the benefit of heterogeneous data," *Pattern Recognition*, vol. 74, pp. 90–109, 2018.

[9] T. Sattler, B. Leibe, and L. Kobbelt, "Efficient & effective prioritized matching for large-scale image-based localization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 9, pp. 1744–1756, 2017.

[10] Y. Kim, J. Jeong, and A. Kim, "Stereo camera localization in 3d lidar maps," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 1–9.

[11] R. W. Wolcott and R. M. Eustice, "Visual localization within lidar maps for automated urban driving," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 176–183.

[12] A. D. Stewart and P. Newman, "Laps-localisation using appearance of prior structure: 6-dof monocular camera localisation using prior pointclouds," in *2012 IEEE International Conference on Robotics and Automation*, 2012, pp. 2625–2632.

[13] G. Pascoe, W. Maddern, and P. Newman, "Direct visual localisation and calibration for road vehicles in changing city environments," in *IEEE International Conference on Computer Vision Workshops*, 2015, pp. 9–16.

[14] P. Neubert, S. Schubert, and P. Protzel, "Sampling-based methods for visual navigation in 3d maps by synthesizing depth images," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 2492–2498.

[15] A. Gawel, T. Cieslewski, R. Dubé, M. Bosse, R. Siegwart, and J. Nieto, "Structure-based vision-laser matching," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 182–188.

[16] D. J. Campbell, L. Petersson, L. Kneip, and H. Li, "Globally-optimal inlier set maximisation for camera pose and correspondence estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.

[17] P. David, D. DeMenthon, R. Duraiswami, and H. Samet, "Simultaneous pose and correspondence determination using line features," in *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2, 2003, pp. II–II.

[18] F. Moreno-Noguer, V. Lepetit, and P. Fua, "Pose priors for simultaneously solving alignment and correspondence," in *European Conference on Computer Vision*, 2008, pp. 405–418.

[19] J. Levinson and S. Thrun, "Automatic online calibration of cameras and lasers." in *Robotics: Science and Systems*, vol. 2, 2013, p. 7.

[20] L. Zhou, Z. Li, and M. Kaess, "Automatic extrinsic calibration of a camera and a 3d lidar using line and plane correspondences," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 5562–5569.

[21] D. Cattaneo, M. Vaghi, A. L. Ballardini, S. Fontana, D. G. Sorrenti, and W. Burgard, "Cmrnet: Camera to lidar-map registration," in *2019 IEEE Intelligent Transportation Systems Conference*, 2019, pp. 1283–1289.

[22] M. Feng, S. Hu, M. H. Ang, and G. H. Lee, "2d3d-matchnet: learning to match keypoints across 2d image and 3d point cloud," in *2019 International Conference on Robotics and Automation*, 2019, pp. 4790–4796.

[23] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Epnp: An accurate o (n) solution to the pnp problem," *International Journal of Computer Vision*, vol. 81, no. 2, p. 155, 2009.

[24] X. Lu, Y. Liu, and K. Li, "Fast 3d line segment detection from unorganized point cloud," *arXiv preprint arXiv:1901.02532*, 2019.

[25] R. G. Von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall, "Lsd: A fast line segment detector with a false detection control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 4, pp. 722–732, 2008.

[26] N. Xue, S. Bai, F. Wang, G.-S. Xia, T. Wu, and L. Zhang, "Learning attraction field representation for robust line segment detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1595–1603.

[27] S. Agarwal, K. Mierle, and Others, "Ceres solver," http://ceres-solver.org.

[28] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The euroc micro aerial vehicle datasets," *The International Journal of Robotics Research*, vol. 35, no. 10, pp. 1157–1163, 2016.

[29] D. Girardeau-Montaut, "Cloudcompare-open source project," *Open-Source Project*, 2011.

[30] K. Huang, Y. Wang, Z. Zhou, T. Ding, S. Gao, and Y. Ma, "Learning to parse wireframes in images of man-made environments," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 626–635.

[31] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 573–580.

[32] Z. Zhang and D. Scaramuzza, "A tutorial on quantitative trajectory evaluation for visual (-inertial) odometry," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 7244–7251.