

LDL: Line Distance Functions for Panoramic Localization

Junho Kim¹, Changwoon Choi¹, Hojun Jang¹, and Young Min Kim^{1,2}

¹ Dept. of Electrical and Computer Engineering, Seoul National University

² Interdisciplinary Program in Artificial Intelligence and INMC, Seoul National University

82magnolia@snu.ac.kr, changwoon.choi00@gmail.com, {j12040208, youngmin.kim}@snu.ac.kr

Abstract

We introduce *LDL*, a fast and robust algorithm that localizes a panorama to a 3D map using line segments. *LDL* focuses on the sparse structural information of lines in the scene, which is robust to illumination changes and can potentially enable efficient computation. While previous line-based localization approaches tend to sacrifice accuracy or computation time, our method effectively observes the holistic distribution of lines within panoramic images and 3D maps. Specifically, *LDL* matches the distribution of lines with 2D and 3D line distance functions, which are further decomposed along principal directions of lines to increase the expressiveness. The distance functions provide coarse pose estimates by comparing the distributional information, where the poses are further optimized using conventional local feature matching. As our pipeline solely leverages line geometry and local features, it does not require costly additional training of line-specific features or correspondence matching. Nevertheless, our method demonstrates robust performance on challenging scenarios including object layout changes, illumination shifts, and large-scale scenes, while exhibiting fast pose search terminating within a matter of milliseconds. We thus expect our method to serve as a practical solution for line-based localization, and complement the well-established point-based paradigm. The code for *LDL* is available through the following link: <https://github.com/82magnolia/panoramic-localization>.

1. Introduction

Estimating the location of a mobile device or agent with respect to a 3D map, widely referred to as visual localization, has vast applications in robotics and AR/VR. Compared to perspective images, which are more widely used for localization, panorama images provide a 360° field of view that contains ample visual evidence from the holistic scene context. In this light, there have been recent advances in visual localization using panoramic images [6, 7, 23, 24]

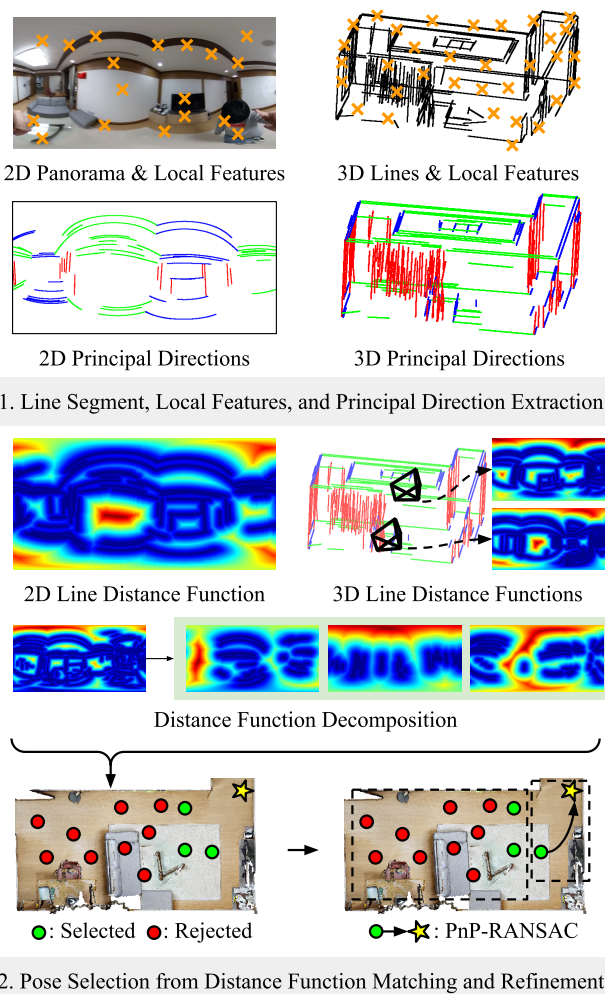


Figure 1. Overview of our approach. *LDL* assumes a 3D map equipped with lines and local features, and similarly preprocesses the 2D panorama prior to localization. *LDL* then selects candidate poses by matching 2D, 3D line distance functions through decomposition along principal directions that effectively represent the sparse geometry of lines. Finally, the selected poses are refined via local feature matching [39] and PnP-RANSAC [13, 25].

that demonstrate reasonably stable localization, with state-of-the-art methods leveraging a two-step process of candidate pose selection and refinement [24, 38]. Nevertheless, many existing methods for this task have limitations in computational efficiency and robustness, mainly stemming from the costly or unstable pose selection process. As global feature descriptors [2, 20] or a large number of colored points [23, 24] are the main components for this step, the pipelines can be memory and compute intensive or fragile to large illumination changes [23, 24].

To overcome such limitations, we explore the alternative direction of using *lines* as the major cue for panoramic localization. Lines have a number of desirable properties compared to commonly used raw color, semantic labels or learned global features [7, 23, 38]. First, due to the long-standing work in line segment extraction [16, 17, 50, 53], it is cheap and stable to extract line segments even amidst dramatic changes in illumination or moderate motion blur. Second, lines are sparse representations of a scene and can potentially lead to small memory consumption and computation. Nevertheless, line segments alone are visually ambiguous compared to other localization cues (color, global features, etc.), which makes them harder to tailor for successful localization. While there exist prior works in line-based visual localization [14, 29, 51], many focus on using lines for *pose refinement* after finding coarse poses from conventional global feature comparisons [14, 51] or exhibit unstable performance compared to conventional point-based methods [29]. Further, prior works often involve expensive line-specific feature extraction to distinguish contexts and establish one-to-one line correspondences [51].

LDL is a fast and robust localization method that leverages the holistic context from lines in panoramas and 3D maps to effectively find the camera pose. In contrast to previous works [14, 51], we retain our focus on using line segments for *pose search* based on the observation that conventional point-based matching [10, 39] performs stably once given a good initial pose. As shown in Figure 1, given a panoramic image of an unknown location, we utilize the distribution of extracted line segments and compare it against those in the pre-captured 3D map. First, the candidate pose selection step rapidly evaluates an immense set of poses within a matter of milliseconds and selects the coarse poses to further optimize. Here LDL compares the distribution of lines in 2D and 3D evaluated on their spherical projections using distance functions, as shown in Figure 1. The distance function imbues relative spatial context even in featureless regions and quickly matches poses without establishing explicit correspondences between detected lines. We further enhance the discriminative power of distance functions by decomposition, and separately evaluate lines aligned with each principal directions. Once a small set of initial poses are found, LDL refines them with PnP-

RANSAC [13, 25], where we leverage powerful local features from recent works [10, 39] to establish good 2D-3D correspondences.

We evaluate LDL in various indoor scenes where it performs competitively against all tested baselines while demonstrating robust performance in scenes with object changes or large illumination shifts. Further, LDL exhibits an order-of-magnitude faster runtime compared to global feature comparison [2, 15, 20] due to the efficient formulation. By only using the geometric information of lines and pre-trained visual features, we expect LDL to serve as a practical localization algorithm that could enhance and complement existing visual localization techniques.

2. Related Work

Line-Based Localization Inspired by abundant straight-lines and rectangular structures in man-made objects, many works attempt visual localization with line segments [1, 14, 29, 47, 51, 52]. Micusik et al. [29] utilize the line segments extracted from the 3D model to directly match line segments in images by comparing the Chamfer distance in 2D and 3D. However, lines, even when perfectly matched, are inherently subject to ambiguity along the line direction. Yoon et al. [51] suggest removing such ambiguities by treating points on a line segment as verbal tokens in natural language processing, where line features are learned using Transformers [48]. Such learning-based approaches are trained with a database of pose-annotated images or require additional computation [14, 51, 52]. Further, these approaches only use lines for pose refinement, assuming a coarse pose estimate to be given via global feature comparisons [2, 15]. LDL takes a different approach and focuses on robust *pose selection* based on lines. We compare LDL against existing approaches for line-based localization, where LDL performs competitively against these methods while balancing robustness and efficiency.

Point-Based Localization Most visual localization algorithms follow a point-based paradigm, focusing on sparse feature point correspondences [8, 21, 26, 32, 36–38, 40–43], dense matching via coordinate regression of scene points [4, 26], or minimizing color discrepancies of dense 3D points via gradient descent [23, 24]. Conventional approaches using a perspective camera input take a two-step approach, where coarse poses are first estimated using global feature descriptors [2, 15] and refined with PnP-RANSAC from local feature matches [10, 27, 39] or dense matches from scene coordinate regression [26, 43]. Recent panoramic localization methods [6, 7, 23, 24] also follow a similar two-step approach, where exemplary methods find candidate poses via color distribution matching and refine them using gradient descent optimization [23, 24]. While these algorithms can robustly handle a modest range of scene changes due to

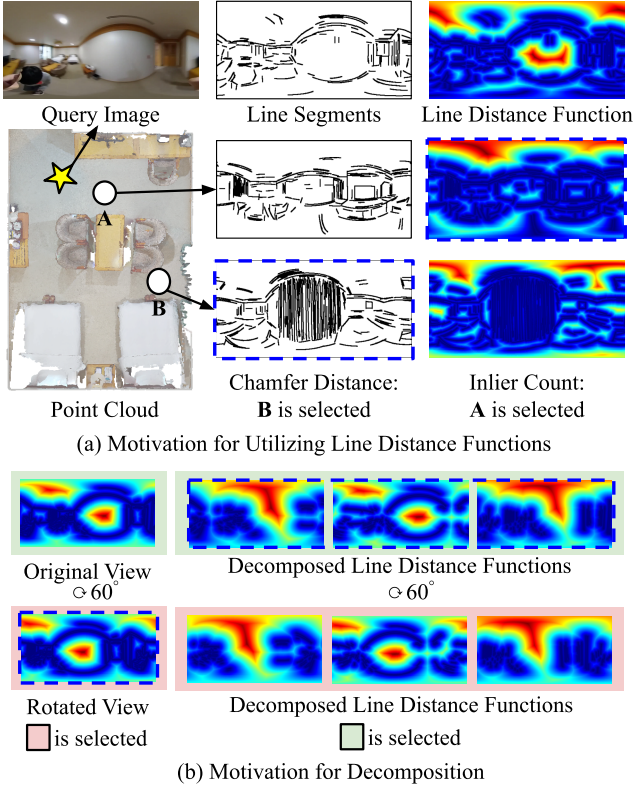


Figure 2. Motivation for (a) utilizing and (b) decomposing line distance functions. (a) Line distance functions disambiguate regions with dense lines. Given two candidate poses close (A) and far (B) from ground truth, Chamfer distance falsely favors B near dense lines, whereas distance functions correctly rank the poses. (b) Decomposition further reduces ambiguities from rotation by separately considering line segments with varying directions. Given an original view close to the ground truth (green) and a rotated view (red), the decomposition better distinguishes the two views by correctly selecting the original view over the rotated view.

the holistic view from panoramas, the algorithms can still fail with significant changes in illumination. We compare LDL against exemplary point-based methods and demonstrate that line segments could be effectively utilized for accurate and robust localization even without the costly calculation of global features or color matching.

3. Method

LDL aims at finding the pose at which the query image I is taken with respect to a 3D scene, where Figure 1 depicts the localization steps taken by LDL. We first represent the 3D scene using a line map equipped with local feature descriptors for keypoint locations, and similarly acquire line segments and local descriptors for the query image prior to localization (Section 3.1). We then estimate the three principal directions for 2D and 3D by voting, from which we can deduce a set of rotations considering the sign and permuta-

tion ambiguity (Section 3.2). Given the fixed set of candidate rotations, we construct an initial set of possible poses incorporating translations. We generate the decomposed line distance functions at each pose and choose the promising poses by comparing the distance functions with a robust loss function (Section 3.3). As the final step, the selected poses are refined by performing PnP-RANSAC [13] using feature matches [39] with the query image (Section 3.4).

3.1. Localization Input Preparation

Map Building LDL operates using a 3D map consisting of line segments and local features. We build such a map starting from a colored point cloud $P = \{X, C\}$. To obtain the 3D line segments we use the line extraction method from Xiaohu et al. [49], which can quickly process point clouds containing millions of points within a few seconds. We further remove short, noisy line segments from the raw detection with a simple filtering step: given the point cloud bounding box of size $b_x \times b_y \times b_z$, we filter out 3D line segments shorter than $\lambda(b_x + b_y + b_z)/3$ with $\lambda = 0.1$ in all our experiments. The 2D line segments are then filtered with an adaptive length threshold to match the filtering rate of 3D line segments. Specifically, we choose the threshold value such that the ratio of lines filtered in 2D equals that in 3D.

To obtain local features embedded in the 3D map, we first render synthetic views at various locations using the point cloud color values. Specifically, we project the input point cloud $P = \{X, C\}$ at a virtual camera and assign the measured color $Y(u, v) = C_n$ at the projected location of the corresponding 3D coordinate $(u, v) = \Pi(RX_n + t)$ to create the synthetic view Y . We then extract local features for each synthetic view Y using SuperPoint [10], and back-project the local features to their 3D locations, which in turn results in keypoint descriptors embedded in 3D space. Note that while we illustrate map building using a colored point cloud, our setup can also work with line-based SfM maps [28, 34, 35] since the input to our pipeline is lines and associated local features.

Panorama Pre-processing Similar to map building, we extract line segments and local features from the query panorama image. We use LSD [16] to acquire line segments, which is a robust line detection algorithm that can stably extract lines even under motion blur or lighting changes. To remove noisy line detections as in the 3D case, we filter 2D line segments with an adaptive length threshold to match the filtering rate of 3D line segments. Specifically, for each scene we choose the threshold value such that the ratio of lines filtered in 2D equals that in 3D. Then, we extract local feature descriptors using SuperPoint [10], where the results will later be used for pose refinement in Section 3.4.

3.2. Candidate Rotation Estimation

Given the detected line segments, LDL first estimates a set of feasible rotations by extracting principal directions, which we define as the most common line directions in 2D and 3D. Let $L_{2D} = \{l\}$ denote the line segments in 2D, where $l = (s, e)$ is a tuple of start point $s \in \mathbb{S}^2$ and end point $e \in \mathbb{S}^2$. Note that we operate on the spherical projection space and treat lines and points on panoramas as arcs and points on the unit sphere \mathbb{S}^2 respectively. Similarly, let $L_{3D} = \{\tilde{l}\}$ denote the line segments in 3D, with $\tilde{l} = (\tilde{s}, \tilde{e})$ being a tuple containing start and end points $\tilde{s}, \tilde{e} \in \mathbb{R}^3$.

LDL estimates the vanishing point and votes for the principal directions in 2D and 3D. In 2D we first extract vanishing points by finding the points of intersection of extended 2D line segments. The 2D principal directions $P_{2D} = \{p\}$ are defined as the top k_{2D} vanishing points containing the most incident lines, where $p \in \mathbb{R}^3$ is a unit norm vector denoting the vanishing point location in the sphere. Similarly, the 3D principal directions $P_{3D} = \{\tilde{p}\}$ are defined as the top k_{3D} most common line directions from 3D line segments obtained via voting. Note that the 3D direction $\tilde{p} \in \mathbb{R}^3$ is also normalized.

LDL estimates the feasible candidate rotations up to uncertainty in the combinatorial ambiguities when matching the principal directions in 2D and 3D. Specifically, we select triplets of directions from P_{2D} and P_{3D} , yielding a total of $\binom{k_{2D}}{3} \times \binom{k_{3D}}{3} \times 3! \times 2^3$ possible combinations, additionally considering the sign and permutation ambiguity. For each pair of triplets, we apply the Kabsch algorithm [22] to find the optimal rotation that aligns the 2D directions to 3D directions. Discarding infeasible rotations that have large mean squared error, we obtain N_r rotations. The possible rotations are further filtered using line distance function presented in the next section.

3.3. Line Distance Functions for Pose Selection

We propose line distance functions to efficiently evaluate a large pool of poses and select promising candidate poses. The initial pool of poses is the combination of possible translations with the rotations found in the previous section. To this end, N_t translations are chosen within grid partitions of the 3D point cloud, where details are explained in the supplementary material. The resulting $N_t \times N_r$ poses are ranked using line distance functions.

Distance Function Definition Distance functions are designed to compare the holistic spatial context captured from the large field of view in panorama images. They are defined for every point including void regions without any lines and can quickly rank poses. Compared to Chamfer distance or learned line embeddings used in prior work [29, 51], LDL does not attempt pairwise matching

between lines, which is often costly and can incur failure modes. For example, it is ambiguous to correctly match between densely packed lines as shown in Figure 2a.

A line distance function is a dense field of distance values to detect lines in the 2D query image or the spherical projection at an arbitrary pose in 3D. For a point x on the unit sphere \mathbb{S}^2 , the 2D line distance function is given as

$$f_{2D}(x; L_{2D}) = \min_{l \in L_{2D}} D(x, l). \quad (1)$$

Here $D(x, l)$ is the spherical distance from x to line segment $l = (s, e)$, namely

$$D(x, l) = \begin{cases} \sin^{-1} \left| \langle x, \frac{s \times e}{\|s \times e\|} \rangle \right| & \text{if } x \in \mathcal{Q}(s, e) \\ \min(\cos^{-1} \langle x, e \rangle, \cos^{-1} \langle x, s \rangle) & \text{otherwise,} \end{cases} \quad (2)$$

where $\mathcal{Q}(s, e)$ is the spherical quadrilateral formed from $\{s, e, \pm(s \times e)/\|s \times e\|\}$.

Similarly, the 3D line distance function is defined for each candidate rotation $R \in SO(3)$ and translation $t \in \mathbb{R}^3$. Using the spherical projection function $\Pi(\cdot) : \mathbb{R}^3 \rightarrow \mathbb{S}^2$ that maps a point in 3D to a point on the unit sphere, the 3D line segment $\tilde{l} = (\tilde{s}, \tilde{e})$ is projected to 2D under the candidate transformation as $l = (\Pi(R\tilde{s} + t), \Pi(R\tilde{e} + t))$. For simplicity, let $\Pi_L(\tilde{l}; R, t)$ denote the projection of a line segment in 3D to the spherical surface. Then the 3D line distance function is defined as follows,

$$f_{3D}(x; L_{3D}, R, t) = \min_{\tilde{l} \in L_{3D}} D(x, \Pi_L(\tilde{l}; R, t)). \quad (3)$$

As shown in Figure 3, one can expect poses closer to the ground truth to have similar 2D and 3D line distance functions. Therefore, we evaluate $N_t \times N_r$ poses according to the similarity of line distance functions.

We apply a robust loss function that measures inlier counts to quantify the affinity of the line distance functions. For each candidate pose $\{R, t\}$ we count the number of points whose distance function differs below a threshold τ ,

$$L(R, t) = -\sum_{q \in Q} \mathbb{1}\{|f_{2D}(q; L_{2D}) - f_{3D}(q; L_{3D}, R, t)| < \tau\}, \quad (4)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function and $Q \subset \mathbb{S}^2$ is a set of query points uniformly sampled from a sphere. The loss function only considers inlier counts, and thus is robust to outliers from scene changes or line misdetections. We validate the efficacy of the robust loss function in Section 4.2.

Distance Function Decomposition To further enhance pose search using line distance functions, we propose to decompose the distance functions along three principal directions. While line distance functions provide useful evidence for line-based localization, they lack a sense of direction as

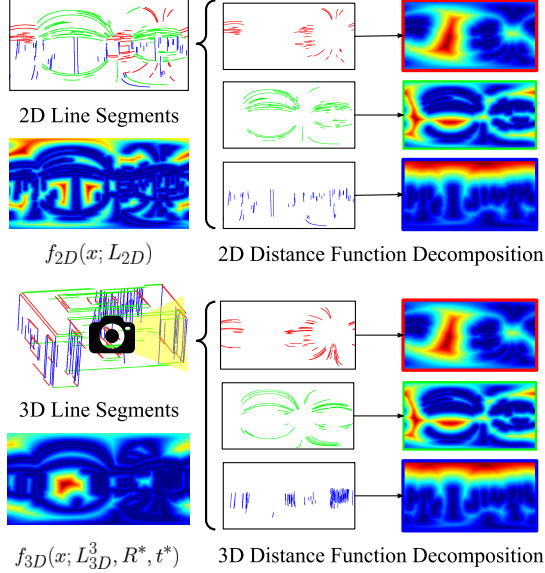


Figure 3. Line distance function visualization and decomposition at the ground truth pose R^*, t^* . LDL decomposes distance functions using principal directions and enhances their expressiveness.

in Figure 2b, where the distance functions alone cannot effectively distinguish rotated views at a fixed translation.

We split line segments along the principal directions used for rotation estimation and define separate line distance functions for each group of lines, as shown in Figure 3. Recall from Section 3.2 that each candidate rotation R is obtained from a pair of triplets in 2D and 3D principal directions denoted as $\hat{P}_{2D}^R = \{p_1, p_2, p_3\}$ and $\hat{P}_{3D}^R = \{\tilde{p}_1, \tilde{p}_2, \tilde{p}_3\}$. We associate line segments that are parallel to directions in $\hat{P}_{2D}^R, \hat{P}_{3D}^R$, leading to three groups of line segments $L_{2D}^R = \{L_{2D}^1, L_{2D}^2, L_{2D}^3\}$ and $L_{3D}^R = \{L_{3D}^1, L_{3D}^2, L_{3D}^3\}$ in 2D and 3D, respectively. We separately define line distance functions for the three groups using Equation 2, namely $f_{2D}(x; L_{2D}^i)$ and $f_{3D}(x; L_{3D}^i, R, t)$ for $i = 1, 2, 3$. Then the robust loss function in Equation 4 can be modified to accommodate the decomposed distance functions,

$$L(R, t) = -\sum_{i=1}^3 \sum_{q \in Q} \mathbb{1}\{|f_{2D}(q; L_{2D}^i) - f_{3D}(q; L_{3D}^i, R, t)| < \tau\}. \quad (5)$$

We validate the importance of distance function decomposition in Section 4.2.

3.4. Candidate Pose Refinement

After we select the top K poses from the pool of $N_t \times N_r$ poses with the loss function values from Equation 5, we refine them using local feature matching as shown in Figure 1. Here we utilize the cached local features from Section 3.1. Specifically, for each selected pose we first retrieve the set of visible 3D keypoints at that pose and perform local fea-

ture matching against the 2D keypoints in the query image. In this process we use SuperGlue [39] for feature matching and select the candidate pose with the most matches. Finally, we apply PnP-RANSAC [13, 19, 25] on the matched 2D and 3D keypoint coordinates to obtain a refined pose estimate. Backed by local feature matching that stably operates given decent coarse estimates from line distance functions, LDL can robustly function as an effective localization method which we further verify in Section 4.

4. Experiments

We evaluate LDL in various localization scenarios and analyze its performance. Our method is mainly implemented using PyTorch [31], and is accelerated with a single RTX 2080 GPU. In all our experiments we set the number of principal directions as $k_{2D}=20, k_{3D}=3$, the inlier threshold $\tau=0.1$, and the number of query points as $|Q|=42$. We report the full hyperparameter setup in the supplementary material. Following prior works [6, 7, 23], we report the median translation and rotation errors along with the localization accuracy where a prediction is considered correct if the translation error is below 0.1m and the rotation error is below 5° .

Datasets We evaluate LDL in two indoor localization datasets: Stanford 2D-3D-S [3] and OmniScenes [23]. Stanford-2D-3D-S [3] contains 1413 panorama images from 272 rooms subdivided into six areas. Each area has diverse indoor scenes such as offices, hallways, and auditoriums where repetitive structure and featureless regions are present. OmniScenes contains 4121 panorama images from seven 3D scans, where the panorama images are captured with cameras either handheld or robot mounted, and at different times of day including large changes in furniture configurations. The dataset has three splits (Robot, Handheld, Extreme) that are recorded in scenes with/without changes, where images in the Extreme split are captured under large camera motion.

Baselines We compare LDL against three point-based baselines (PICCOLO, CPO, structure-based) and two line-based baselines (Chamfer distance-based, Line Transformer [51]). PICCOLO (PC) [23] and the follow-up work CPO [24] is an optimization-based algorithm that finds pose by minimizing the color discrepancy between the point cloud and the query image. Structure-based approach [38, 46] (SB) is one of the most prominent methods for visual localization using perspective cameras. We implement a method for panorama images, where candidate poses are retrieved from an image database using a global feature extractor [15] and further refined using SuperGlue [39] matches. For fair comparison, we undistort the

Dataset	t -error (m)						R -error ($^{\circ}$)						Accuracy					
	PC	SB	CD	LT	CPO	LDL	PC	SB	CD	LT	CPO	LDL	PC	SB	CD	LT	CPO	LDL
Area 1	0.02	0.02	0.12	0.02	0.01	0.02	0.46	0.62	1.14	0.62	0.25	0.54	0.66	0.89	0.50	0.90	0.90	0.86
Area 2	0.76	0.04	1.16	0.04	0.01	0.02	2.25	0.72	11.54	0.72	0.27	0.66	0.45	0.76	0.35	0.74	0.81	0.77
Area 3	0.02	0.03	0.79	0.02	0.01	0.02	0.49	0.57	4.54	0.55	0.24	0.54	0.57	0.92	0.36	0.88	0.78	0.89
Area 4	0.18	0.02	0.33	0.02	0.01	0.02	4.17	0.57	1.97	0.56	0.28	0.48	0.49	0.91	0.46	0.91	0.83	0.88
Area 5	0.50	0.03	0.95	0.03	0.01	0.02	14.64	0.69	41.84	0.65	0.27	0.54	0.44	0.80	0.36	0.79	0.74	0.81
Area 6	0.01	0.02	0.50	0.02	0.01	0.02	0.31	0.63	1.20	0.60	0.18	0.50	0.69	0.88	0.47	0.87	0.90	0.83
Total	0.03	0.03	0.73	0.02	0.01	0.02	0.63	0.63	2.30	0.63	0.24	0.53	0.54	0.85	0.39	0.84	0.83	0.83

Table 1. Localization performance evaluation in Stanford 2D-3D-S [3], compared against PICCOLO (PC) [23], structure-based approach (SB), Chamfer distance-based approach (CD), Line Transformer (LT) [51], and CPO [24].

Dataset	Accuracy					
	PC	SB	CD	LT	CPO	LDL
Original	0.45	0.69	0.21	0.68	0.72	0.89
Gamma	0.00	0.63	0.47	0.59	0.00	0.82
Intensity	0.00	0.56	0.40	0.58	0.80	0.76
White Balance	0.00	0.62	0.32	0.67	0.74	0.91

Table 2. Localization accuracy on synthetic color variations applied to Room 3 in the Extreme split from OmniScenes [23].

panorama image into cubemaps and perform feature matching, where the results are then fed to PnP-RANSAC for refinement. In addition, we construct the database of pose-annotated images by rendering synthetic views at various locations in the colored point cloud.

Chamfer distance-based approach (CD), inspired from Micusik et al. [29], ranks candidate poses by comparing the spherical Chamfer distance of line segments in the synthetic views against the query image. Line Transformer by Yoon et al. [51] (LT) ranks candidate poses using Transformer-based [48] matching learned for each line segment. As this baseline also requires a pose-annotated database, we construct a synthetic database similar to the structure-based approach, and apply the undistortion process for fair comparison. We provide additional details about the baselines in the supplementary material.

4.1. Localization Evaluation

Stanford 2D-3D-S We first assess the localization performance of LDL against the baselines in the Stanford 2D-3D-S dataset, as shown in Table 1. LDL performs competitively against the strong baselines (Structure-based and Line Transformer) that apply powerful neural networks for candidate pose search. While the dataset contains hallways and auditoriums with large featureless regions or repetitive structure, LDL leverages the holistic distribution of lines using distance functions and shows stable performance without resorting to costly neural network computations. Further, LDL shows superior performance when compared against the Chamfer distance-based method, which indicates that solely focusing on line matches for ranking candidate poses can lead to suboptimal performance.



Figure 4. Color variations for evaluating illumination robustness.

OmniScenes We additionally compare LDL against baselines in the OmniScenes dataset, as shown in Table 3. Unlike the Stanford 2D-3D-S dataset, all images exhibit blur from camera motion and approximately half of the images contain changes in object layout. In splits not containing changes, LDL performs competitively against the baselines, which supports our claim that line distance functions enable effective pose search without using neural networks. Further, LDL attains the highest accuracy in splits containing scene changes and notably in the extreme split that contains the largest amount of motion blur. This is due to the stable line extraction [16, 17, 50, 53] that enables resilience against motion blur, and the robust distance function comparison (Equation 4) that rejects outliers for handling scene changes. We further verify the importance of each components in LDL in Section 4.2.

Illumination Robustness Evaluation To validate the illumination robustness of LDL, we measure localization performance after applying synthetic color variations. We select Room 3 from the Extreme split in OmniScenes for evaluation. As shown in Figure 4, the image gamma, white balance, and average intensity are modified to an arbitrary value, where further details are deferred to the supplementary material. We report the results of LDL along with the baselines in Table 2. CPO, PICCOLO, and the structure-based baseline all suffer from performance degradation, as the color values are directly utilized for finding initial poses. Notably, Yoon et al. [51] also shows performance drop, as Transformer line features are affected by the illumination changes of the image. As LDL relies on the spatial structure of line segments for candidate pose search, it is robust to illumination variations, leading to stable performance across all color variations. Further, note that while all the methods excluding PICCOLO [23] and CPO [24] use local feature

Split	Change	t -error (m)						R -error ($^{\circ}$)						Accuracy					
		PC	SB	CD	LT	CPO	LDL	PC	SB	CD	LT	CPO	LDL	PC	SB	CD	LT	CPO	LDL
Robot	✗	0.02	0.03	1.74	0.03	0.01	0.02	0.27	0.58	89.23	0.59	0.12	0.49	0.69	0.99	0.31	0.99	0.89	0.98
Hand	✗	0.01	0.03	2.10	0.03	0.01	0.03	0.23	0.63	89.02	0.64	0.13	0.54	0.81	0.95	0.29	0.95	0.80	0.97
Robot	✓	1.07	0.04	1.78	0.04	0.02	0.03	21.03	0.64	89.27	0.65	1.46	0.58	0.41	0.93	0.30	0.94	0.59	0.95
Hand	✓	0.53	0.04	1.70	0.04	0.02	0.03	7.54	0.71	88.50	0.70	0.37	0.64	0.47	0.92	0.30	0.90	0.60	0.92
Extreme	✓	1.24	0.04	1.55	0.04	0.03	0.03	23.71	0.83	88.54	0.84	0.37	0.72	0.41	0.89	0.29	0.88	0.59	0.92

Table 3. Localization performance evaluation in OmniScenes [23], considering both scenes with and without object layout changes.

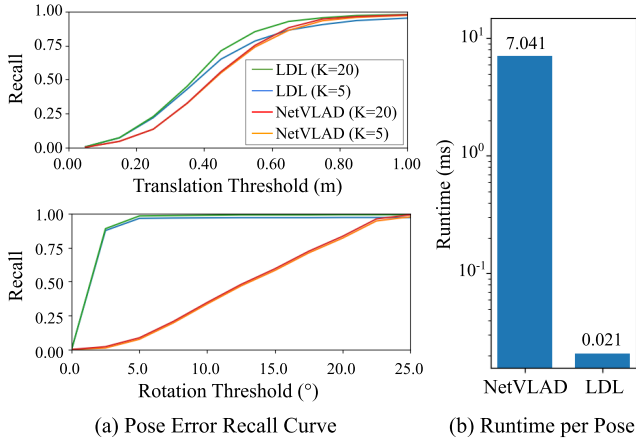


Figure 5. Pose error recall and runtime comparison between candidate pose search using LDL and NetVLAD [2].

Method	t -error (m)	R -error ($^{\circ}$)	Acc.	Component	CPU	GPU
SB ($K=10$)	0.06	1.18	0.63	Line Segment Extraction	0.141	0.141
SB ($K=20$)	0.05	1.07	0.71	Rotation Estimation	1.124	0.009
LDL ($K=10$)	0.07	1.36	0.63	Distance Function Computation	0.052	0.001
LDL ($K=20$)	0.07	1.27	0.69	Candidate Pose Refinement	5.573	0.587
				Total Runtime (sec)	6.890	0.738

(a) Multi-Room Localization

(b) Runtime on CPU and GPU

Table 4. Multi-room localization compared against Structure-Based method (SB) with various number of candidate poses (K) and runtime analysis of LDL.

matching for pose refinement, there is a large performance gap between LDL and the other methods. This validates our focus on designing a stable candidate pose selection method, as modern feature descriptors and matching algorithms [10, 11, 38, 39] are fairly robust against adversaries such as illumination changes.

4.2. Performance Analysis

Candidate Pose Search Evaluation To evaluate the efficacy of line distance functions for candidate pose search, we compare the retrieval accuracy of LDL against NetVLAD [2], which is a widely used global feature extractor [20, 38, 51]. Note that NetVLAD is used as the candidate pose selection module in the structure-based baseline. We use the Extreme split from OmniScenes for evaluation, where the translation and rotation error recall curve along

with the runtime for processing a single candidate pose is reported in Figure 5. For fair comparison we use the identical pool of translations for both methods as $N_t = 50$ and assign a large number of candidate rotations for NetVLAD with $N_r = 216$. Additional setup details are reported in the supplementary material. While neural network-based pose search methods can perform city-scale search [2, 15, 18], the line distance functions in LDL exhibit competitive performance to NetVLAD in indoor environments. The distance functions provide highly discriminative spatial context, which enables effective pose search. Furthermore, the runtime for pose search in LDL is much shorter than NetVLAD, due to the highly efficient computation of distance functions only conducted on sparse sphere points. This is in contrast to NetVLAD where visual features are computed with a neural network for each view. The line distance functions enable quick and effective pose initialization, which in turn allow LDL to be usable in various practical localization scenarios.

Runtime Analysis We analyze the runtime of LDL in Table 4b where we decompose the runtime for localizing a single query image from OmniScenes [23]. We assume that 3D scanning along with map building is done offline and only consider the computation time for online operations, namely 2D line segment extraction, candidate pose selection and refinement. Overall, the pose selection process including rotation estimation and distance function computation exhibits a small runtime for both CPU and GPU, which validates the efficiency of our proposed line-based pose search. Nevertheless, the pose refinement exhibits a relatively larger runtime, which is mainly due to the large number of features in panoramas compared to normal images with a smaller field of view. While we attained our focus in pose search and used the off-the-shelf local feature matching algorithms for pose refinement [10, 39], devising highly efficient feature matching algorithms tailored specifically for panoramas is left as future work.

Scalability Analysis We assess the scalability of LDL to large-scale indoor scenes using the OmniScenes [23] dataset. While the previous set of experiments assume room-scale localization scenarios, here we test LDL us-

ing the entire OmniScenes dataset as the 3D map. Table 4a shows the localization results, where LDL is compared against the structure-based method at various number of candidate poses (K). LDL exhibits performance on a par with the structure-based method, which indicates that line distance functions can scalably handle large scenes consisting of multiple rooms. Nevertheless, scaling LDL to even larger scale scenes (e.g. building-scale scenes as in InLoc [46]) is left as future work.

Privacy Preservation Analysis While the main goal of LDL is to offer fast and robust localization based on lines, we find that with a small modification our method can offer light-weight privacy protection in client-server localization scenarios [5, 9, 12, 44, 45]. Following prior works [30, 45], we consider the case where a client using an edge device wants to localize oneself against a 3D line map stored in the cloud. Privacy breaches occur if the service provider maliciously tries to view the visual data captured by the client. This is possible even when only the local feature descriptors are shared between the client and server, by using feature inversion methods [33] that reconstruct the original image from a sparse set of local features as shown in Figure 6.

By changing LDL to only exploit local features near lines during refinement, we can prevent privacy breaches including feature inversion attacks without largely sacrificing localization performance. First, as LDL uses line segments for candidate pose selection the clients only need to share the extracted line segments with the service providers for initial pose search, instead of the entire view that would be needed for global feature-based methods. Second, as local features near line segments are shared with the service provider for pose refinement, feature inversion methods cannot faithfully recover the original visual content. We validate this claim with a small set of experiments performed in the Stanford 2D-3D-S dataset [3], where we filter descriptors whose spherical distances to the nearest line segment are over 0.05 rad. As shown in Figure 6, this line-based filtering degrades the quality of feature inversion attacks by hiding objects that potentially contain sensitive information while only incurring small drops in localization accuracy. We report additional details and results regarding the potential of LDL for privacy preservation in the supplementary material.

4.3. Ablation Study

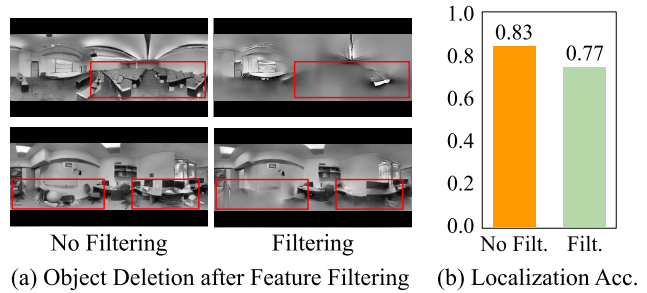
We ablate the distance function decomposition, number of query points, and robust loss function, which are key components of LDL in the OmniScenes Extreme split. In Table 5a, LDL is first compared against the baseline that does not apply decomposition and use the loss function in Equation 4. Decomposition leads to a large performance gain, as the distance functions are further disambiguated

Method	t -error (m)	R -error (°)	Acc.	Method	t -error (m)	R -error (°)	Acc.
w/o Decomposition	1.00	3.97	0.37	w/ L1 Loss	0.08	1.38	0.55
w/ $ Q = 10$	0.04	0.85	0.77	w/ L2 Loss	0.17	1.48	0.34
w/ $ Q = 21$	0.04	0.71	0.88	w/ Huber Loss	0.11	1.39	0.50
w/ $ Q = 84$	0.03	0.66	0.95	w/ Median Loss	0.08	1.22	0.55
Ours ($ Q = 42$)	0.03	0.72	0.92	Ours	0.07	1.22	0.68

(a) Decomposition & Query Points

(b) Choice of Loss Function

Table 5. Ablation study of various components of LDL.



(a) Object Deletion after Feature Filtering

(b) Localization Acc.

Figure 6. Visualization of feature inversion attacks on panoramic inputs along with the localization accuracy before and after line-based feature filtering

and split into each principal direction. We further test the effect of the number of query points $|Q|$ on evaluating the robust loss function. While increasing the number of query points enhances performance, the improvement is not as significant and incurs additional computation. Conversely, using a smaller number of query points lead to ambiguities in distance function matching, exhibiting poor performance. The number of query points $|Q| = 42$ balances both the computational efficiency and localization accuracy of LDL. We finally validate the robust loss function in Equation 5 by comparing LDL against variants using other loss functions: L1, L2, Huber, and Median loss. Here we report results from the Wedding Hall scene, as this scene contains drastic scene changes with large amounts of outliers. As shown in Table 5b, inlier counting proposed in Equation 5 attenuates outliers and exhibits optimal performance, demonstrating the effectiveness of the robust loss function.

5. Conclusion

We presented LDL, a fast and robust algorithm for panorama to point cloud localization using line segments. LDL benefits from the illumination-robustness of line segments and the holistic context of panoramas by using a novel formulation based on line distance functions. The distance functions effectively handle visual ambiguities of line segments, as they provide spatial meaning to void regions often neglected by existing line-based localization methods. In addition, by evaluating distance functions only on sparsely sampled query points, LDL performs rapid candi-

date pose search with accuracy on a par with learning-based global feature extractors. As a result, LDL performs robust localization in various challenging scenarios with a short runtime. We expect LDL to complement and enhance the currently prevalent point-based localization algorithms for highly robust and practical localization.

Acknowledgements This work was partially supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2023-00218601), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2021-0-02068, Artificial Intelligence Innovation Hub), and Samsung Electronics Co., Ltd. Young Min Kim is the corresponding author.

References

- [1] Hichem Abdellali, Robert Frohlich, Viktor Vilagos, and Zoltan Kato. L2d2: Learnable line detector and descriptor. In *2021 International Conference on 3D Vision (3DV)*, pages 442–452, 2021. [2](#)
- [2] R. Arandjelović, P. Gronat, A. Torii, T. Pajdla, and J. Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. [2](#), [7](#)
- [3] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. [5](#), [6](#), [8](#)
- [4] Eric Brachmann, Alexander Krull, Sebastian Nowozin, J. Shotton, F. Michel, S. Gumhold, and C. Rother. Dsac — differentiable ransac for camera localization. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2492–2500, 2017. [2](#)
- [5] Christian Cachin, Idit Keidar, and Alexander Shraer. Trusting the cloud. *SIGACT News*, 40(2):81–86, jun 2009. [8](#)
- [6] Dylan Campbell, Lars Petersson, Laurent Kneip, and Hongdong Li. Globally-optimal inlier set maximisation for camera pose and correspondence estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, page preprint, June 2018. [1](#), [2](#), [5](#)
- [7] Dylan Campbell, Lars Petersson, Laurent Kneip, Hongdong Li, and Stephen Gould. The alignment of the spheres: Globally-optimal spherical mixture alignment for camera pose estimation. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, page to appear, Long Beach, USA, June 2019. IEEE. [1](#), [2](#), [5](#)
- [8] Gabriela Csurka and Martin Humenberger. From handcrafted to deep local invariant features. *CoRR*, abs/1807.10254, 2018. [2](#)
- [9] Deeksha Dangwal, Vincent T. Lee, Hyo Jin Kim, Tianwei Shen, Meghan Cowan, Rajvi Shah, Caroline Trippel, Brandon Reagen, Timothy Sherwood, Vasileios Balntas, Armin Alaghi, and Eddy Ilg. Analysis and mitigations of reverse engineering attacks on local feature descriptors, 2021. [8](#)
- [10] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *CVPR Deep Learning for Visual SLAM Workshop*, 2018. [2](#), [3](#), [7](#)
- [11] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-Net: A Trainable CNN for Joint Detection and Description of Local Features. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. [7](#)
- [12] Mihai Dusmanu, Johannes Schönberger, Sudepta Sinha, and Marc Pollefeys. Privacy-preserving image features via adversarial affine subspace embeddings. In *Computer Vision and Pattern Recognition (CVPR 2021)*. CVF/IEEE, June 2021. [8](#)
- [13] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981. [1](#), [2](#), [3](#), [5](#)
- [14] Shuang Gao, Jixiang Wan, Yishan Ping, Xudong Zhang, Shuzhou Dong, Jijunnan Li, and Yandong Guo. Pose refinement with joint optimization of visual points and lines. *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2888–2894, 2021. [2](#)
- [15] Yixiao Ge, Haibo Wang, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-supervising fine-grained region similarities for large-scale image localization. In *European Conference on Computer Vision*, 2020. [2](#), [5](#), [7](#)
- [16] Rafael Grompone von Gioi, Jeremie Jakubowicz, Jean-Michel Morel, and Gregory Randall. Lsd: A fast line segment detector with a false detection control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(4):722–732, 2010. [2](#), [3](#), [6](#)
- [17] Geonmo Gu, Byungsoo Ko, SeoungHyun Go, Sung-Hyun Lee, Jingeun Lee, and Minchul Shin. Towards real-time and light-weight line segment detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022. [2](#), [6](#)
- [18] Stephen Hausler, Sourav Garg, Ming Xu, Michael Milford, and Tobias Fischer. Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14141–14152, 2021. [7](#)
- [19] J. A. Hesch and S. I. Roumeliotis. A direct least-squares (dls) method for pnp. In *2011 International Conference on Computer Vision*, pages 383–390, 2011. [5](#)
- [20] Martin Humenberger, Yohann Cabon, Nicolas Guerin, Julien Morat, Jérôme Revaud, Philippe Rerole, Noé Pion, Cesar de Souza, Vincent Leroy, and Gabriela Csurka. Robust image retrieval-based visual localization using kapture, 2020. [2](#), [7](#)
- [21] Martin Humenberger, Yohann Cabon, Nicolas Guérin, Julien Morat, Jérôme Revaud, Philippe Rerole, Noé Pion, César Roberto de Souza, Vincent Leroy, and Gabriela Csurka. Robust image retrieval-based visual localization using kapture. *CoRR*, abs/2007.13867, 2020. [2](#)
- [22] W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A*, 32(5):922–923, Sep 1976. [4](#)

- [23] Junho Kim, Changwoon Choi, Hojun Jang, and Young Min Kim. Piccolo: Point cloud-centric omnidirectional localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3313–3323, October 2021. [1](#), [2](#), [5](#), [6](#), [7](#)
- [24] Junho Kim, Hojun Jang, Changwoon Choi, and Young Min Kim. Cpo: Change robust panorama to point cloud localization. *ECCV*, 2022. [1](#), [2](#), [5](#), [6](#)
- [25] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Eppn: An accurate o(n) solution to the pnp problem. *International Journal Of Computer Vision*, 81:155–166, 2009. [1](#), [2](#), [5](#)
- [26] Xiaotian Li, Shuzhe Wang, Yi Zhao, Jakob Verbeek, and Juho Kannala. Hierarchical scene coordinate classification and regression for visual localization. In *CVPR*, 2020. [2](#)
- [27] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60:91–110, 2004. [2](#)
- [28] Branislav Micusik and Horst Wildenauer. Structure from motion with line segments under relaxed endpoint constraints. In *2014 2nd International Conference on 3D Vision*, volume 1, pages 13–19, 2014. [3](#)
- [29] Branislav Micusik and Horst Wildenauer. Descriptor free visual indoor localization with line segments. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3165–3173, 2015. [2](#), [4](#), [6](#)
- [30] Tony Ng, Hyo Jin Kim, Vincent T. Lee, Daniel DeTone, Tsun-Yi Yang, Tianwei Shen, Eddy Ilg, Vassileios Balntas, Krystian Mikolajczyk, and Chris Sweeney. Ninjadesc: Content-concealing visual descriptors via adversarial learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12797–12807, June 2022. [8](#)
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. [5](#)
- [32] N. Pion, M. Humenberger, G. Csurka, Y. Cabon, and T. Sattler. Benchmarking image retrieval for visual localization. In *2020 International Conference on 3D Vision (3DV)*, pages 483–494, Los Alamitos, CA, USA, nov 2020. IEEE Computer Society. [2](#)
- [33] Francesco Pittaluga, Sanjeev J Koppal, Sing Bing Kang, and Sudepta N Sinha. Revealing scenes by inverting structure from motion reconstructions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 145–154, 2019. [8](#)
- [34] Yohann Salaün, Renaud Marlet, and Pascal Monasse. Robust and accurate line- and/or point-based pose estimation without manhattan assumptions. In *European Conference on Computer Vision*, 2016. [3](#)
- [35] Y. Salaun, R. Marlet, and P. Monasse. Line-based robust sfm with little image overlap. In *2017 International Conference on 3D Vision (3DV)*, pages 195–204, Los Alamitos, CA, USA, oct 2017. IEEE Computer Society. [3](#)
- [36] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. *CoRR*, abs/1812.03506, 2018. [2](#)
- [37] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. *CoRR*, abs/1911.11763, 2019. [2](#)
- [38] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *CVPR*, 2019. [2](#), [5](#), [7](#)
- [39] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *CVPR*, 2020. [1](#), [2](#), [3](#), [5](#), [7](#)
- [40] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Improving image-based localization by active correspondence search. In Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid, editors, *Computer Vision – ECCV 2012*, pages 752–765, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. [2](#)
- [41] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Efficient & effective prioritized matching for large-scale image-based localization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(9):1744–1756, 2017. [2](#)
- [42] Johannes L. Schonberger, Hans Hardmeier, Torsten Sattler, and Marc Pollefeys. Comparative evaluation of hand-crafted and learned local features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. [2](#)
- [43] J. Shotton, B. Glocker, C. Zach, S. Izadi, A. Criminisi, and A. Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2930–2937, 2013. [2](#)
- [44] Pablo Speciale, Johannes L. Schönberger, Sing Bing Kang, Sudepta N. Sinha, and Marc Pollefeys. Privacy preserving image-based localization. 2019. [8](#)
- [45] Pablo Speciale, Johannes Schönberger, Sudepta Sinha, and Marc Pollefeys. Privacy preserving image queries for camera localization. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1486–1496, 2019. [8](#)
- [46] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. InLoc: Indoor Visual Localization with Dense Matching and View Synthesis. In *CVPR 2018 - IEEE Conference on Computer Vision and Pattern Recognition*, Salt Lake City, United States, June 2018. [5](#), [8](#)
- [47] Felix Taubner, Florian Tschoop, Tonci Novkovic, Roland Siegwart, and Fadri Furrer. Lcd – line clustering and description for place recognition, 2020. [2](#)
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon,

- U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [2](#), [6](#)
- [49] Lu Xiaohu, Liu Yahui, and Li Kai. Fast 3d line segment detection from unorganized point cloud. *arXiv preprint arXiv:1901.02532*, 2019. [3](#)
- [50] Nan Xue, Song Bai, Fudong Wang, Gui-Song Xia, Tianfu Wu, and Liangpei Zhang. Learning attraction field representation for robust line segment detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#), [6](#)
- [51] Sungho Yoon and Ayoung Kim. Line as a visual sentence: Context-aware line descriptor for visual localization. *IEEE Robotics and Automation Letters*, 6(4):8726–8733, 2021. [2](#), [4](#), [5](#), [6](#), [7](#)
- [52] Xin Yu, Sagar Chaturvedi, Chen Feng, Yuichi Taguchi, Teng-Yok Lee, Clinton Fernandes, and Srikumar Ramalingam. Vlase: Vehicle localization by aggregating semantic edges. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3196–3203, 2018. [2](#)
- [53] Ziheng Zhang, Zhengxin Li, Ning Bi, Jia Zheng, Jinlei Wang, Kun Huang, Weixin Luo, Yanyu Xu, and Shenghua Gao. Ppgnet: Learning point-pair graph for line segment detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#), [6](#)