

# How Privacy-Preserving are Line Clouds? Recovering Scene Details from 3D Lines

Kunal Chelani<sup>1</sup> Fredrik Kahl<sup>1</sup> Torsten Sattler<sup>1,2</sup>

<sup>1</sup>Chalmers University of Technology <sup>2</sup>Czech Technical University in Prague

## Abstract

*Visual localization is the problem of estimating the camera pose of a given image with respect to a known scene. Visual localization algorithms are a fundamental building block in advanced computer vision applications, including Mixed and Virtual Reality systems. Many algorithms used in practice represent the scene through a Structure-from-Motion (SfM) point cloud and use 2D-3D matches between a query image and the 3D points for camera pose estimation. As recently shown, image details can be accurately recovered from SfM point clouds by translating renderings of the sparse point clouds to images. To address the resulting potential privacy risks for user-generated content, it was recently proposed to lift point clouds to line clouds by replacing 3D points by randomly oriented 3D lines passing through these points. The resulting representation is unintelligible to humans and effectively prevents point cloud-to-image translation. This paper shows that a significant amount of information about the 3D scene geometry is preserved in these line clouds, allowing us to (approximately) recover the 3D point positions and thus to (approximately) recover image content. Our approach is based on the observation that the closest points between lines can yield a good approximation to the original 3D points. Code is available at <https://github.com/kunalchelani/Line2Point>.*

## 1. Introduction

Visual localization is the problem of estimating the position and orientation from which an image was taken in a known scene. Visual localization is a fundamental part of computer vision systems such as self-driving cars [29, 66], Augmented and Mixed Reality applications [2, 12], Structure-from-Motion (SfM) [28, 31, 65, 69], and Simultaneous Localization and Mapping (SLAM) [18, 51, 57].

Classical approaches to visual localization [16, 30, 40–42, 56, 59, 62, 86, 87] are based on local features such as SIFT [43]. They use SfM to construct a sparse 3D point cloud of the scene, where each point is associated with the local image features it was triangulated from. Descriptor matching between local features extracted in a test / query

image and the 3D points then yields a set of 2D-3D matches that can be used for RANSAC-based camera pose estimation [9, 17, 23, 36–39].

Traditionally, work on visual localization has focused on accurate and scalable algorithms able to cover large areas [15, 30, 40, 58, 74, 76, 86] or to run in real-time on mobile devices with limited memory and compute capabilities [2, 42, 46, 47, 50]. Thus, the underlying scene representations have been designed to enable efficient 2D-3D matching [30, 47, 58, 61, 63] and / or to limit memory consumption [10, 11, 41, 46, 47, 61]. Privacy aspects such as avoiding user generated content from being recovered either through 3D models stored in the cloud or through query images sent to a server have traditionally not been taken into account.

Recently, [55, 70] showed that it is possible to recover images from SfM point clouds. Given a rendering of the point cloud (and the feature descriptors associated with the 3D points), [55] uses a CNN to translate the rendering to a complete image. Their work clearly demonstrates that storing SfM point clouds creates potential privacy risks as an attacker could recover details from user-uploaded content stored in the cloud. To prevent such attacks, [71] proposed to replace each SfM point through a random line passing through this point (*cf.* Fig. 1(left)). They showed that the resulting representation is unintelligible to humans, prevents a direct application of [55], and still enables accurate camera pose estimation. This idea of lifting points to lines was later adapted for privacy-perserving SLAM [67]. However, this paper shows that it is possible to (approximately) recover the original 3D point positions from a line cloud (*cf.* Fig. 1(middle)), again enabling us to use [55, 70] to obtain images (*cf.* Fig. 1(right)).

In detail, this paper makes the following **contributions**: (i) in the case that the line directions are chosen uniformly at random, as is the case in [71], we show that knowledge about local neighborhoods allows us to (approximately) recover the original 3D points based on the closest points between pairs of lines. (ii) based on this insight, we propose a two-stage approach that first recovers these neighborhoods and then estimates the 3D points corresponding to the input lines. (iii) detailed experiments on both indoor and outdoor

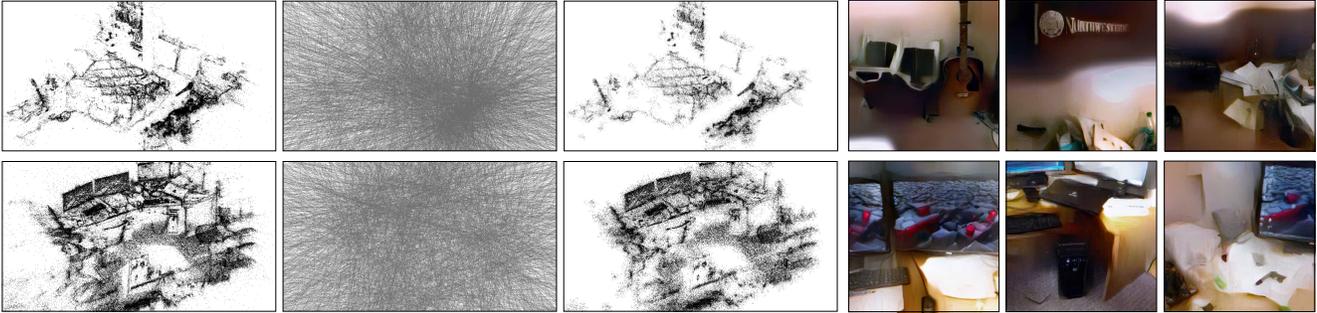


Figure 1. In order to preserve privacy, [71] proposed to store line clouds instead of point clouds for visual localization (left). While unintelligible to the human eye, we show that it is possible to recover the underlying 3D point clouds (middle). Applying a point cloud-to-image translation approach [55] then allows us to recover image details (right), showing that lifting point clouds to line clouds can still preserve privacy critical information that can later be extracted from the line clouds.

datasets show that our approach allows us to faithfully reproduce the original point clouds. In addition, applying [55] on the resulting point clouds enables us to recover image details (*cf.* Fig. 1(right)). (iv) while using line clouds alone is not effective in obfuscating the underlying 3D scene geometry, we show that using (very) sparse line clouds effectively prevents our approach from recovering image details. Our results clearly show that further research on privacy-preserving scene representations is needed.

## 2. Related Work

Privacy-preserving methods in computer vision are not new. For instance, it is standard to blur faces and license plates to prevent the identification of persons and cars, respectively. Yet, it has been shown that this process is not privacy preserving [53] as one can still train a recognition system on obfuscated images with reasonable accuracy rates. Privacy preserving image degradation with adversarial training has been proposed for visual recognition [83]. In this paper, we focus on privacy preserving localization.

**Visual localization.** Traditionally, visual localization approaches have relied on local features such as SIFT [43] or its learned alternatives [3, 22, 54]. Structure-based approaches [16, 30, 40, 41, 56, 60, 61, 74, 86] represent a scene through a SfM point cloud, with each point being associated with at least one local image feature. 2D-3D matches between features extracted from a new image and the 3D points in the SfM model can then be used for pose estimation [9, 23, 26, 27, 36, 37]. State-of-the-art approaches for long-term localization, *i.e.*, robust camera pose estimation under changing conditions, follow this approach, but use learned features [22, 25, 58, 75]. An alternative to using a 3D point cloud is to either compute the position and orientation of a test image from relative poses to database images [87, 88], or to compute a 3D model on-the-fly [62]. Still, these approaches rely on local image features and are thus susceptible to the image inversion approach from [55].

Learned localization approaches either replace the complete localization pipeline [8, 33, 34, 81] or the 2D-3D

matching stage [4–7, 13, 14, 19, 49, 68] through machine learning. The former, which directly regress a camera pose via a convolutional neural network (CNN), have been shown to perform similar to image retrieval techniques [64], *i.e.*, nearest neighbor classifiers that only approximate the pose of the test image [1, 77, 78]. The second family of learned localization methods, which regress a 3D scene coordinate for each pixel in a test image, has been shown to achieve high pose accuracy on small scenes [5–7, 13, 14, 68]. However, they currently do not scale well to larger or more complex scenes [6, 13]. As such, such learned localization systems are currently not used in practice.

**Recovering image content from features.** The descriptors of local features represent an abstract representation of a patch centered around a keypoint. A long-standing result is that it is possible to recover the image content from gradient-based features such as SIFT [82] and HOG [80], even if the descriptors are quantized into visual words [32]. Naturally, training CNNs to recover images improves the reconstruction quality [20, 21]. Further work shows that the deep representations learned by neural networks can be inverted to recover images, which can be used as a tool to visualize what such networks learn [48, 84, 85].

**Privacy-preserving visual localization.** Pittaluga *et al.* extended these results on recovering images from 2D local features to 3D Structure-from-Motion point clouds [55]. They showed that a CNN can be trained to recover an image from the projection of 3D points (together with their descriptors) into a synthetic view. They concluded that scene representations based on SfM point clouds can allow an attacker to recover private details. Recently, [70] showed that high-quality images can also be recovered from sparse colored point clouds without image descriptors. In order to enable privacy preserving localization, Speciale *et al.* proposed to replace each SfM point by a random line passing through that point [71], where the direction of the line is sampled uniformly at random from a unit sphere. They showed that the resulting line clouds

still enable precise camera pose estimation. They argued that line clouds preserve privacy as they prevent the approaches from [55, 70] from being applicable, which would ensure that user-recorded scenes can be safely stored in the cloud. Follow-up work to this seminal paper showed how to adapt SLAM systems to integrate this idea into a SLAM system [67] and how to enable privacy-preserving image queries for localization [72] and privacy-preserving SfM [24]. [24, 72] operate on 2D rather than 3D representations and replace each 2D image feature by a 2D line.

This paper investigates the claim of preserving user privacy by lifting 3D point to line clouds made by [71]. We show that it is possible to (approximately) recover the underlying point positions using only the provided line geometry. As a result, we show that image-level details can be obtained from line clouds through recovering the underlying 3D point cloud. Our approach is based on the observation that two random 3D lines often enough have their closest points nearby to the original 3D points. However, this does not hold in 2D, *i.e.*, our approach is only applicable to recover 3D point clouds but not 2D feature clouds.

### 3. From Point Clouds to Line Clouds

Structure-based visual localization approaches [30, 40, 41, 60, 86] use 2D-3D correspondences between pixels and 3D points in a scene model to estimate the camera pose of a given query image. To this end, classical feature-based approaches represent the scene through a 3D point cloud  $\mathcal{P} = \{(\mathbf{p}_i, \mathbf{d}_i)\}$ , where each 3D point  $\mathbf{p}_i \in \mathbb{R}^3$  is associated with one or more image feature descriptors  $\mathbf{d}_i$ , *e.g.*, a 128-dimensional SIFT [43] descriptor.<sup>1</sup> Pittaluga *et al.* showed that it is possible to “invert” the point clouds, often constructed using SfM, used by feature-based localization systems [55]. More precisely, they showed that it is possible to use a CNN to recover image details from a projection of a sparse set of 3D point and their descriptors into an image. They concluded that the commonly used point cloud scene representations do not preserve privacy.

To avoid revealing details of a user-uploaded scene model through the inversion process, Speciale *et al.* [71] propose to lift the underlying point cloud  $\mathcal{P} = \{(\mathbf{p}_i, \mathbf{d}_i)\}$  to a line cloud  $\mathcal{L} = \{(\mathbf{l}_i, \mathbf{d}_i)\}$ . Each point  $\mathbf{p}_i$  is replaced by a random line  $\mathbf{l}_i$  passing through it.<sup>2</sup> This introduces an additional degree of freedom, namely the true position of the point  $\mathbf{p}_i$  along the line  $\mathbf{l}_i$  is unknown. Since these 3D lines project to lines in 2D, the inversion approach from [55] is not directly applicable anymore. Furthermore, the resulting scene representation is unintelligible to humans (*cf.*

<sup>1</sup> $\mathcal{P} = \{(\mathbf{p}_i, \mathbf{d}_i)\}$  is the minimally required scene representation. Some methods such as [40, 60] store additional details such as co-visibility information. We only use the minimal representation in this paper.

<sup>2</sup>The chosen line representation, *e.g.*, a Plücker vector as in [71], is not important in the context of this paper.

Fig. 1). Consequently, Speciale *et al.* claim that the “3D line cloud representation hides the underlying scene geometry and prevents the extraction of sensitive information” [71]. Yet, Sec. 4 shows that lifting a point cloud to a line cloud does not completely hide the underlying geometric properties if the line directions are uniformly sampled from a unit sphere. Based on this insight, we develop an algorithm to recover point clouds from line clouds.

## 4. Recovering Point Clouds from Line Clouds

Considered in isolation, a single line  $\mathbf{l}$  is perfectly privacy-preserving as all points on  $\mathbf{l}$  are equally good candidates for the true 3D point  $\mathbf{p}$  that gave rise to the line. However, not all points along the line will be equally likely when taking other lines into account. This is due to the fact that the points in the original point cloud are not randomly distributed but lie on surfaces. Information about the distribution might be preserved in a line cloud. In this paper, we show that it can be possible to recover information about the local neighborhood of points from a line cloud. In turn, this information can be used to recover points from lines.

A line cloud is characterized by the distribution of the directions of lines drawn through the underlying points. This design choice can be used to choose a distribution that conceals the most information. Speciale *et al.* propose to sample line directions independently and uniformly at random for each point, as this distribution helps to ensure good localization accuracy [71]. Consequently, we assume that the line directions are drawn independently from a uniform distribution over a unit sphere and denote the resulting line clouds as *uniform line clouds*. Sec. 4.1 shows that this distribution implies that the two closest points between two lines  $\mathbf{l}_i, \mathbf{l}_j$  are likely to be relatively close to the original 3D points  $\mathbf{p}_i, \mathbf{p}_j$ . Sec. 4.2 shows how to leverage this information to recover 3D points from 3D lines. Sec. 4.3 then discusses limitations of our approach, including listing conditions under which point recovery might not be possible.

### 4.1. Information in Uniform Line Clouds

If we consider the point and line clouds,  $\mathcal{P}$  and  $\mathcal{L}$ , to be random variables, then the posterior distribution  $P(\mathcal{P}|\mathcal{L})$  can be obtained via Bayes rule:

$$P(\mathcal{P}|\mathcal{L}) = P(\mathcal{L}|\mathcal{P})P(\mathcal{P})/P(\mathcal{L}) \propto P(\mathcal{L}|\mathcal{P})P(\mathcal{P}) . \quad (1)$$

Since all line directions are drawn independently from another, we have  $P(\mathcal{L}|\mathcal{P}) = \prod_{i=1}^N P(\mathbf{l}_i|\mathbf{p}_i)$ . The probability  $P(\mathbf{l}_i|\mathbf{p}_i)$  of a line  $\mathbf{l}_i$  given its corresponding point  $\mathbf{p}_i$  is zero if  $\mathbf{p}_i$  does not lie on  $\mathbf{l}_i$ . Otherwise, the probability is constant as all line directions are equally likely. Thus, the likelihood function  $P(\mathcal{L}|\mathcal{P})$  is piece-wise constant, *i.e.*, two point clouds  $\mathcal{P}, \mathcal{P}'$  will have the same likelihood  $P(\mathcal{L}|\mathcal{P}) = P(\mathcal{L}|\mathcal{P}')$  as long as every point  $\mathbf{p}_i$  lies on

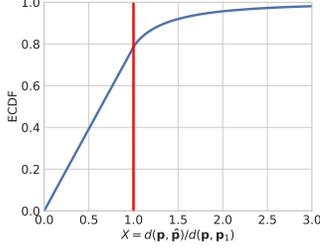


Figure 2. Empirical cumulative distribution of the relation between the Euclidean distance  $d(\mathbf{p}, \mathbf{p}_1)$  of two points  $\mathbf{p}, \mathbf{p}_1$  and the distance  $d(\mathbf{p}, \hat{\mathbf{p}})$  between  $\mathbf{p}$  and an estimate  $\hat{\mathbf{p}}$  for  $\mathbf{p}$  obtained via the closest points of the two lines corresponding to  $\mathbf{p}$  and  $\mathbf{p}_1$ . As can be seen, there is a high chance that  $\mathbf{p}$  is closer to  $\hat{\mathbf{p}}$  than to  $\mathbf{p}_1$ , where the red line denotes equal distance.

its line  $\mathbf{l}_i$ . Consequently, a maximum a posteriori estimate is obtained by maximizing the prior  $P(\mathcal{P})$ , under the constraint that all points should lie on their lines. Unfortunately, defining or learning a general prior distribution  $P(\mathcal{P})$  seems like a hard problem. We thus reason about local neighborhoods instead of the global point cloud. The neighboring points / lines are then used to recover a point position estimate  $\hat{\mathbf{p}}_i$  from the line  $\mathbf{l}_i$ .

Consider two 3D points,  $\mathbf{p}$  and a nearby point  $\mathbf{p}_1$ , as well as their corresponding lines  $\mathbf{l}, \mathbf{l}_1$ . A simple approach to obtain an estimate for  $\mathbf{p}$  is to find the point  $\hat{\mathbf{p}}$  on the line  $\mathbf{l}$  with minimum Euclidean distance to  $\mathbf{l}_1$ . In the following, we show that for uniform line clouds,  $\hat{\mathbf{p}}$  is likely to be a relatively good estimate for  $\mathbf{p}$ . To this end, we study the distribution of the random variable  $X = d(\mathbf{p}, \hat{\mathbf{p}}) / d(\mathbf{p}, \mathbf{p}_1)$ , where  $d(\cdot, \cdot)$  is the Euclidean distance between two 3D points. We empirically measure this distribution by fixing the distance  $d(\mathbf{p}, \mathbf{p}_1) = 1$  and randomly sampling the directions of the two lines  $\mathbf{l}, \mathbf{l}_1$  for one million iterations. Fig. 2 shows the resulting cumulative distribution of  $X$ . A key observation that can be drawn from this is that in nearly 80% of the cases,  $d(\mathbf{p}, \hat{\mathbf{p}})$  is smaller than  $d(\mathbf{p}, \mathbf{p}_1)$ . This implies that if  $d(\mathbf{p}, \mathbf{p}_1)$  is small, there is a good chance that  $\hat{\mathbf{p}}$  will be close to the true point position  $\mathbf{p}$ .

The above analysis can be further extended to a neighborhood of size  $k$ . Let  $\mathcal{N}^k(\mathbf{p}) = \{1, 2, \dots, k\}$  be the indices of the neighboring points of  $\mathbf{p}$ . Let  $\{\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_k\}$  be the lines through these neighboring points and let

$$\mathbf{d}_{\max} = \max_{j \in \mathcal{N}^k(\mathbf{p})} d(\mathbf{p}, \mathbf{p}_j) \quad (2)$$

be the maximum distance between  $\mathbf{p}$  and any of its neighbors. Since all directions are drawn independently, uniformly at random from a unit sphere, the result from the 2-point analysis from above holds pairwise for each pair  $(\mathbf{p}, \mathbf{p}_j)$ ,  $j \in \mathcal{N}^k(\mathbf{p})$ . Thus, from the  $k$  estimates obtained on  $\mathbf{l}$  using lines  $\{\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_k\}$ ,  $0.8k$  can be expected to lie within a distance  $\mathbf{d}_{\max}$  from  $\mathbf{p}$ . If  $\mathbf{d}_{\max}$  is small, this leads to a clustering of estimates close to the true point  $\mathbf{p}$ .

The results from above suggest that given information

about the  $k$  nearest neighbors of each point, it should be possible to obtain accurate point estimates from a line cloud. We verify this intuition through a simple experiment on an indoor scene (*cf.* Fig. 3(a)) for  $k = 50$ . For each point, we obtain 50 estimates on its line using the lines through its nearest neighboring points. Simply taking the median of these estimates produces the result shown in Fig. 3(b). To measure the impact of imperfect neighborhoods, we randomly replace 50% / 90% of the neighbors with random lines from the line cloud. As can be seen in Fig. 3(c), even with 50% outliers, it is still possible to recover the underlying point cloud. However, 90% outliers lead to a very noisy point cloud and the images obtained via [55] from this point cloud become unintelligible (*cf.* Fig. 3(d)). For comparison, Fig. 3(e) shows results obtained with our approach, introduced in Sec. 4.2, that aims to recover the neighborhood from all lines in the line cloud.

## 4.2. Recovering Points from Uniform Line Clouds

As motivated above, estimating a 3D point position  $\mathbf{p}_i$  from a given line  $\mathbf{l}_i$  can be modelled as a two-stage process: the first stage, **neighborhood estimation**, identifies the neighboring 3D points of  $\mathbf{p}_i$  through their corresponding lines. For each such line  $\mathbf{l}_j$ , the point on  $\mathbf{l}_i$  closest to  $\mathbf{l}_j$  provides an estimate for the true point position  $\mathbf{p}_i$ . Given these candidate positions along  $\mathbf{l}_i$ , the second stage then selects a single candidate. This is achieved by finding high-density regions of the candidates along  $\mathbf{l}_i$  via **peak finding**.

Multiple iterations of these two steps are performed to improve the estimates. In the first iteration, only line-to-line distances can be computed and used to select a neighborhood of lines for  $\mathbf{l}_i$ . Starting from the second iteration, point estimates are available for each line. These point estimates can be used to obtain better neighborhood estimates, which in turn lead to better 3D point predictions (*cf.* Fig. 3).

**Neighborhood Estimation.** In the first iteration, we can only measure the minimum distance  $d_{ll}(\mathbf{l}_i, \mathbf{l}_j)$  between lines  $\mathbf{l}_i$  and  $\mathbf{l}_j$ , defined as the Euclidean distance between the closest points on the lines. The minimum distance between two lines cannot be larger than the distance between the corresponding points. Thus, the set  $\mathcal{N}_{ll}^K(\mathbf{l}_i) = \{i_1, \dots, i_K\}$  containing the indices of the  $K$  nearest lines to  $\mathbf{l}_i$  should contain part of the true neighborhood  $\mathcal{N}^k(\mathbf{p}_i)$  if  $K > k$  is large enough.

Once point position estimates  $\{\hat{\mathbf{p}}_i | i \in 1, \dots, N\}$  for the input lines  $\{\mathbf{l}_i | i \in 1, \dots, N\}$  are obtained, we compute two additional neighborhoods:  $\mathcal{N}_{lp}^K(\hat{\mathbf{p}}_i) = \{i_1, \dots, i_K\}$  is the set of indices belonging to the  $K$  lines that have the smallest Euclidean distances to the point estimate  $\hat{\mathbf{p}}_i$ . Similarly,  $\mathcal{N}_{pl}^K(\mathbf{l}_i)$  is the set of indices belonging to the  $K$  point estimates that have the smallest Euclidean distances to the line  $\mathbf{l}_i$ . Intuitively, we expect  $\mathcal{N}_{pl}^K(\mathbf{l}_i)$  to contain those true neighbors that have been estimated close to their true 3D

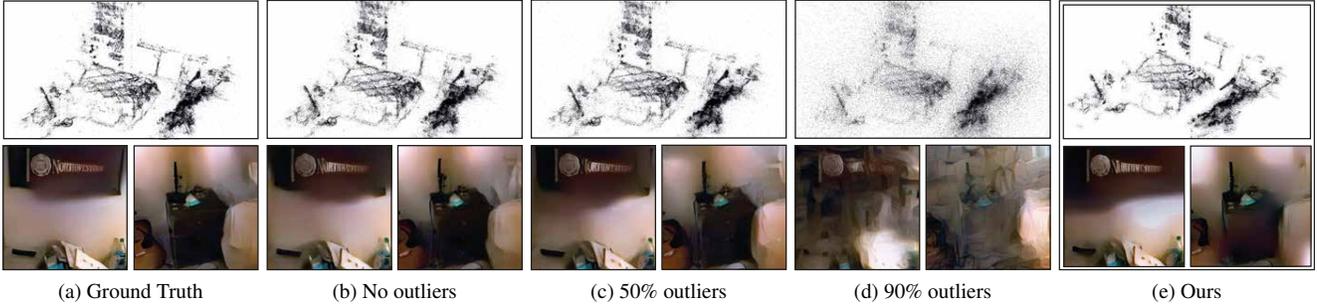


Figure 3. Top row: (a) ground truth point cloud, (b) - (d) point clouds recovered using true neighborhoods with different outlier contamination levels. (e) the result from our method. Bottom row: Images reconstructed from the point clouds using [55].

position in the previous iteration. In practice, we observe that  $\mathcal{N}_{pl}^K$  often has a comparatively high overlap with the true neighborhood of points in  $\mathcal{N}^k(\mathbf{p}_i)$ . However, it still contains outliers corresponding to 3D points estimated in regions through which  $\mathbf{l}_i$  passes. Similarly, we can expect  $\mathcal{N}_{lp}^K(\hat{\mathbf{p}}_i)$  to contain lines corresponding to the true neighboring points from  $\mathcal{N}^k(\mathbf{p}_i)$ , but also outliers from lines from unrelated points that pass through the region containing  $\hat{\mathbf{p}}_i$ . We thus estimate the neighborhood of  $\mathbf{l}_i$  as the intersection  $\mathcal{N}_{pl}^K(\mathbf{l}_i) \cap \mathcal{N}_{lp}^K(\hat{\mathbf{p}}_i)$ .

In our experience, an initial neighborhood size of  $K$  between 300 and 500 for  $\mathcal{N}_{ll}^K(\mathbf{l}_i)$ , followed by a smaller neighborhood size of 100 to 200 for  $\mathcal{N}_{pl}^K(\mathbf{l}_i)$  and  $\mathcal{N}_{lp}^K(\hat{\mathbf{p}}_i)$  works well for a line cloud of about 100k lines.

**Peak Finding.** Let  $\mathcal{N}^K(\mathbf{l}_i)$  be the neighborhood for line  $\mathbf{l}_i$  from the neighborhood estimation stage. Each line  $\mathbf{l}_j$ ,  $j \in \mathcal{N}^K(\mathbf{l}_i)$ , provides a 3D point candidate estimate  $\hat{\mathbf{p}}_{ij}$  for  $\mathbf{p}_i$ , defined as the point on  $\mathbf{l}_j$  with the minimum distance to  $\mathbf{l}_i$ . Let  $\mathbf{o}_i$  be an arbitrary 3D point on  $\mathbf{l}_i$  and  $\mathbf{v}_i$  be the 3D direction of the line. We can then parameterize the candidates as a distribution of scalar values

$$\mathcal{E}_i = \{\beta_{ij} | \hat{\mathbf{p}}_{ij} = \mathbf{o}_i + \beta_{ij}\mathbf{v}_i, \forall j \in \mathcal{N}^K(i)\} \quad (3)$$

along the line. We select the candidate 3D position  $\hat{\mathbf{p}}_i$  for  $\mathbf{p}_i$  by finding high-density regions in this distribution. Following [44, 45], we use the Kuiper’s statistic [35], a non-parametric test statistic that can be used to measure where a given cumulative distribution function (CDF)  $F_X(x)$  differs the most from a reference CDF  $F_T(x)$ , to identify high density regions. We measure the unweighted empirical CDF

$$F_i(x) = \frac{1}{K} \sum_{j=1}^K I_{\beta_{ij} < x} \quad (4)$$

describing the distribution of the points from  $\mathcal{E}_i$  on  $\mathbf{l}_i$ . Here,  $I_{\beta_{ij} < x}$  is an indicator variable taking value 1 if  $\beta_{ij} < x$  and 0 otherwise. We compare  $F_i(x)$  with the CDF  $F_U(x)$  of a uniform distribution between the minimum and maximum value from  $\mathcal{E}_i$ . We compute the two points  $\bar{x}^- = \operatorname{argmax}_x (F_U(x) - F_i(x))$  and  $\bar{x}^+ = \operatorname{argmax}_x (F_i(x) - F_U(x))$  corresponding to the positions

along the line where the two distributions differ most. The differences between the distributions at these points are given as  $D^- = (F_U(\bar{x}^-) - F_i(\bar{x}^-))$  and  $D^+ = (F_i(\bar{x}^+) - F_U(\bar{x}^+))$  and the Kuiper’s statistic is then defined as  $\text{KS} = D^- + D^+$ .

Intuitively,  $\bar{x}^-$  and  $\bar{x}^+$  define the start and end points of a high-density region of the points along the line and KS provides a measure for how much this density differs from a uniform distribution. Recursion within the range  $\bar{x}^-$  and  $\bar{x}^+$  can be used to refine the range. The value of KS can be used to decide when we are sufficiently close to the peak and we stop once KS drops below 0.4. The median value within the range is then used as the estimate  $\hat{\mathbf{p}}_i$  for  $\mathbf{p}_i$ . Similarly, recursion outside the range  $\bar{x}^-$  and  $\bar{x}^+$  can be used to identify multiple candidates and we select the one with the largest KS value. Please see the supp. material for a more detailed description of the peak finding procedure.

### 4.3. Limitations

Our approach fails if it either is not able to recover enough true neighbors in the first stage (*cf.* Fig. 3) or if good position estimates cannot be found in the second stage.

Naturally, changing the distribution of line directions such that the closest points between pairs of lines are far away from the original points will cause both stages to fail. However, the localization approach presented in [71] is based on modelling pose estimation as a relative pose problem for generalized cameras. This allows [71] to estimate the absolute scale of the translation and thus the absolute pose of the test image. In the case that all lines used for pose estimation intersect in the same point, the relative pose problem degenerates to the classical perspective relative problem [52]. In this case, the translation can only be recovered up to a unknown scaling factor and localization fails. As such, there is a trade-off between pose accuracy and preventing the use of closest points on lines. While out of scope for this work, we believe that this is an interesting direction for future research (*cf.* Sec. 6).

An alternative strategy to prevent our method from recovering accurate 3D point positions is to represent the scene as a sparse line cloud. Using only a subset of the orig-

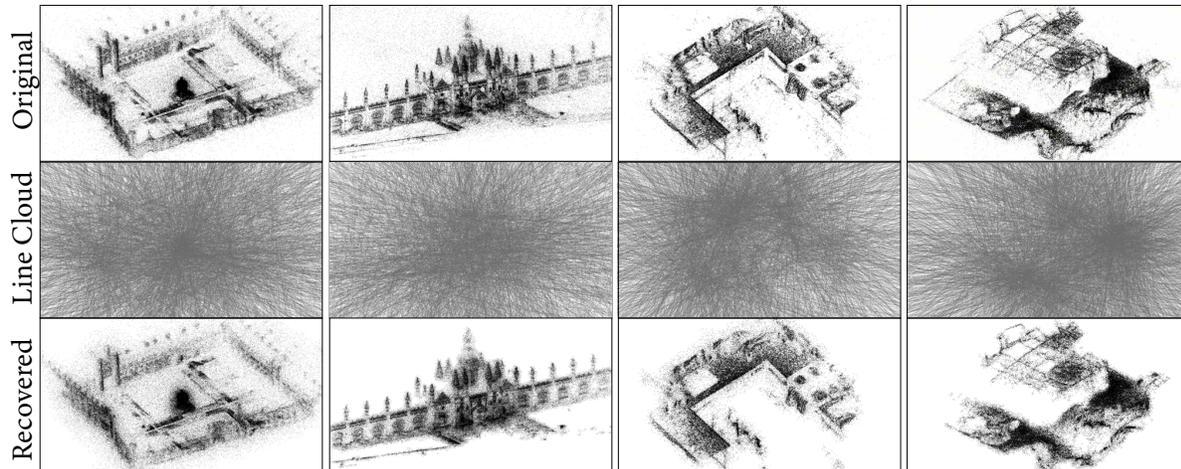


Figure 4. Qualitative results showing point clouds recovered for two scenes each from the outdoor Cambridge Landmark [34] and the indoor 12 Scenes [79] datasets. From left to right, we have the ‘Great Court’, ‘Kings College’, ‘Apt2-Kitchen’ and ‘Apt2-Bedroom’ scenes. For the Cambridge scenes, 30% of all available lines were used for the shown recovered point cloud, while for the indoor scenes from 12 Scenes, all the lines were used. A statistical outlier removal step is performed on the recovered point clouds before visualization.

inal 3D points Consequently, it becomes harder to identify true neighbors based on line-line and point-line distances in the first stage of our algorithm. Even if we can recover the true neighborhood, increasing distances between the true neighbors decreases the chance that the closest points on two lines are close to the true point positions (*cf.* Sec. 4.1). Thus, peak finding might fail if the line clouds are too sparse. Since [71] showed that accurate localization is still possible when using 5-10% of the original lines, we consider this case in our experiments.

Linked to the challenges induced by using sparse line clouds is that our approach struggles to recover points in areas with low point density. Consider the line of a point from a sparse region in the original point cloud that passes through a region containing more points. In this case, there is a good chance that most of the neighbors identified by our method will come from this denser region rather than the original neighborhood. This results in predicting the point position in the wrong part of the scene.

## 5. Experimental Evaluation

In this section, we evaluate our approach for recovering point clouds from line clouds on a set of indoor and outdoor scenes. We first show qualitative and quantitative results for the recovered 3D point clouds. Next, we analyze to what degree image content can be recovered by applying the SfM inversion procedure from [55] on our recovered point clouds. Finally, we show that very sparse line clouds can be used to prevent [55] from obtaining human-interpretable images from our recovered point clouds. Overall, our experiments show that lifting point clouds to line clouds by itself is not sufficient to obfuscate the underlying scene geometry.

**Datasets.** For evaluation, we use the Cambridge Land-

marks [34] and 12 Scenes [79] datasets. The Cambridge Landmarks dataset depicts individual outdoor landmarks. The 12 Scenes dataset depicts 12 smaller rooms in indoor scenes. In addition, we use the Strecha Fountain [73] scene and a dataset depicting castle Dagstuhl in Germany.

**Recovering point clouds.** Fig. 4 shows the point clouds recovered by our method for two outdoor and two indoor scenes. As can be seen, our approach is able to faithfully reproduce the overall structure of the scene. At the same time, our method is also able to reveal details such as the presence of the kitchen sink and the circular burner plates in the kitchen scene, or the chequered design of the bed linen in the bedroom scene. We observe that regions of higher point density are recovered in more detail. This is consistent with the fact that shorter distances between neighboring points increase the chance of accurately recovering the point positions. As predicted in Sec. 4.3, sparser regions are often not well-recovered. Rather, the corresponding point estimates fall into regions with a higher point density. More qualitative results are shown in Fig. 6 and in the supp. material.

To quantify the accuracy with which our method recovers point clouds, we measure the Euclidean distance between the original and recovered point positions. Fig. 5(left) and (middle) show the cumulative distributions of these errors for each of the indoor and outdoor scenes. As can be seen, our approach is able to recover a large fraction of the points within 5cm of their original position for the indoor scenes, with median errors in the range of 1-3cm. In contrast, the errors are considerably larger for the outdoor scenes. This can be explained by the fact that the accuracy with which a point can be recovered depends on the underlying distances to its estimated neighbors. Outdoors, where the structure is farther away from the camera than in indoor scenes, these distances are larger. Furthermore, we only

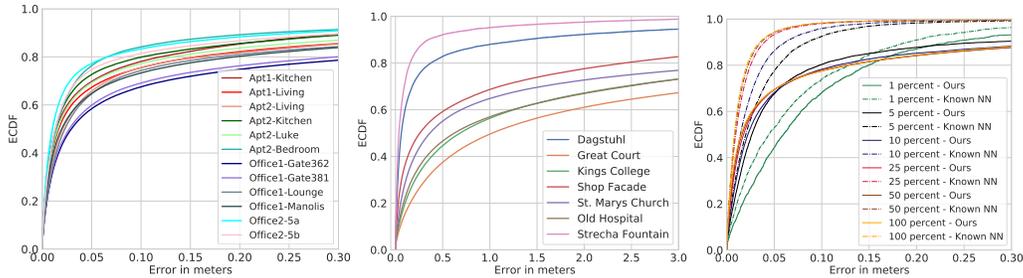


Figure 5. Quantitative results showing the cumulative distribution of errors in the recovering point positions. (Left) Results for each scene from the 12 scenes datasets. (Middle) Results for the outdoor datasets. (Right) Results averaged over four indoor scenes for different sparsity levels. “Known NN” denotes a variant of our method where the true neighborhood is given by an oracle.

use 30% of all lines, selected uniformly at random, for the Cambridge scenes for computational efficiency. Still, Fig. 4 shows that the accuracy of the recovered points is sufficient to produce human-interpretable point clouds.

**Recovering image details.** The main motivation behind using line instead of point clouds was to prevent recovering image content from SfM point clouds [71]. In the next experiment, we thus apply the SfM inversion approach from [55] on our recovered point clouds to see to which degree image content can be recovered from them. We apply the pre-trained model from [55], trained on SfM datasets, without fine-tuning it on our more noisy point clouds.

Fig. 6 shows qualitative results for two outdoor and two indoor scenes, comparing the inversion results of the original and our recovered point clouds. As can be expected, our more noisy point clouds lead to artifacts such as blurry regions and wavy lines instead of straight ones. Still, the obtained images clearly reveal scene information like the overall structure of the buildings, shapes of doors and windows, as well as the presence of small objects such as cushions, shoes, a Bayern Munich flag, *etc.* Based on these results, we conclude that lifting point clouds to line clouds does not guarantee that image details cannot be recovered. Additional qualitative results can be found in the supp. mat.

**(Very) sparse line clouds prevent image recovery.** As discussed in Sec. 4.3, using sparser line clouds still allows accurate localization while potentially preventing our approach from recovering accurate point clouds. In this experiment, following [71], we thus consider sparse line clouds obtained by randomly selecting a subset of the lines.

Fig. 5(right) shows quantitative results averaged over four indoor scenes for different levels of sparsity. We show results for our method and a variant that receives the true neighbors from an oracle. As can be seen, our approach can recover the point positions rather accurately when using down to 5% of the original line cloud. However, there is a significant drop in accuracy when using only 1% of all lines. The comparison with the oracle shows that there is considerable room for improvement in terms of better recovering the true neighborhood of each point / line. However, even

using the oracle still results in a considerable drop in performance when using 1% of all lines.

Fig. 7 shows qualitative results for the images obtained via [55] from the recovered point clouds. While Fig. 5(right) shows that our method provides rather accurate position estimates even at densities as low as 10% and 5%, the resulting images show very little details as compared to images from the original sparse point cloud. Part of the reason is that points from sparser regions are wrongly recovered in denser regions, causing problems for the inversion process. At the same time, sparser point clouds have been shown to be harder to recover in the first place [55]. As can be seen, meaningful image details cannot be recovered when using 5% or less of the lines, even when given the true neighborhood by some oracle. Based on the results, we conclude that sparse line clouds can effectively prevent the recovery of image content using existing methods. Additional qualitative results can be found in the supp. mat.

## 6. Conclusion

In this paper, we have shown that lifting point clouds to line clouds does not necessarily obfuscate the underlying 3D scene structure. We have shown that it is possible to (approximately) recover the underlying 3D point cloud by identifying local neighborhoods. In turn, these neighborhoods are used to obtain point position estimates using closest points between lines, which we have shown to often provide a good approximation to the original 3D points. Quantitative and qualitative results show that our approach enables us to recover image details from line clouds. However, our results also show that sparsification can effectively prevent recovering image details. In the context of privacy-preserving visual localization, we thus conclude that using lines alone does not guarantee privacy-preservation, but that using sparse representations is similarly important. This conclusion is based on the ability of existing SfM inversion methods to recover images from sparse and noisy point clouds. Yet, better recovery algorithms might require us to revisit this problem in the future.

As shown in Fig. 7, using very sparse point clouds is also effective in preventing [55] from recovering image de-

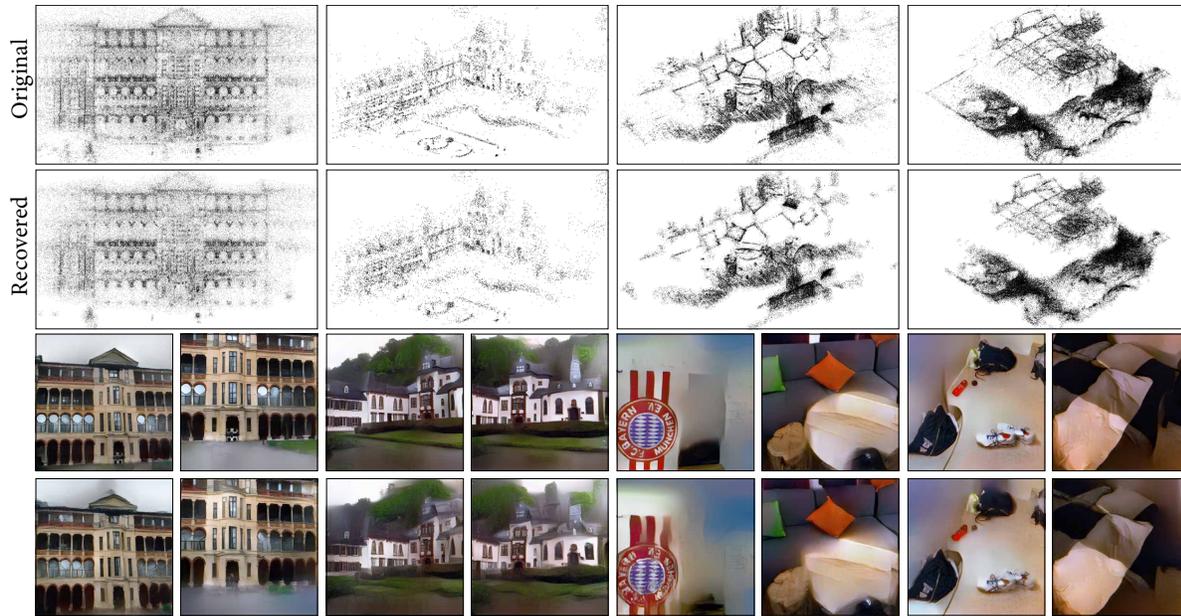


Figure 6. Qualitative results for the recovered point clouds and the images obtained by applying the inversion technique from [55] on the original (above) and recovered (below) point clouds. Left to right: 'Old Hospital', 'Dagstuhl', 'Office1-Lounge', 'Apt2-Bed'.

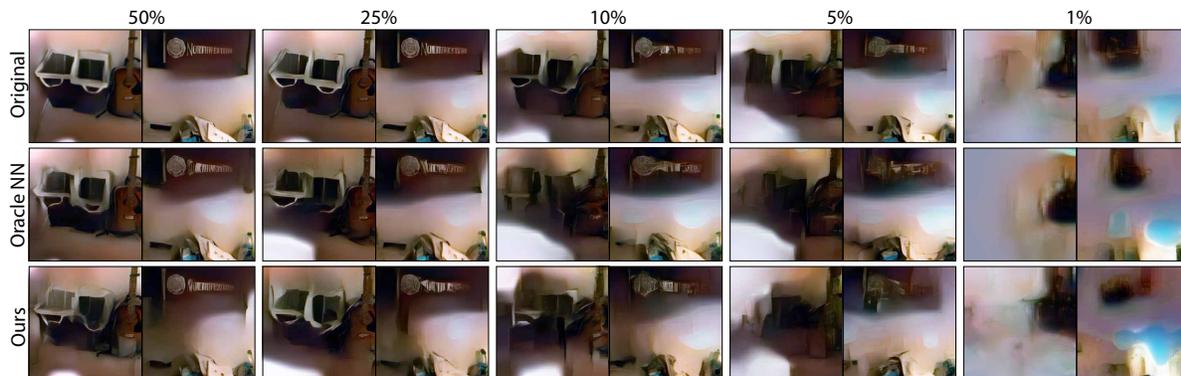


Figure 7. Qualitative results (on scene 'Apt2-Luke' from 12 scenes) for studying the impact of using sparser line clouds to represent the scene. We compare the images obtained by inversion of the original point cloud (top), from a point cloud recovered by our approach using the true neighborhood given by an oracle (middle), and from the point cloud recovered by our method without the oracle (bottom).

tails. Since point-based localization methods are more accurate than line-based ones [71], finding sparsity levels at which point clouds are privacy-preserving while enabling more accurate pose estimates is an interesting direction for future work. Further, our results show room for improvement for our method. Including descriptor information into the recovery process might help unlock this potential.

As detailed in Sec. 4.3, exploring different distributions for line directions is another interesting research direction. Potential approaches to define the distribution include: 1) ensuring that there are no closest points between lines that are within a given threshold of the original points. 2) creating line neighborhoods that do not contain closest points, *e.g.*, by ensuring that lines from many far away points come very close to each other far away from the original point position. 3) ensuring that most lines go through dense neigh-

borhoods, thus allowing to recover these parts while preventing recovery in all the sparser parts. Computing such distribution likely cannot be done individually per point anymore but rather requires complex iterative schemes. At the same time, iteratively applying our approach with a human in-the-loop might allow us to still handle such distributions: after visual inspection, the human removes lines from the neighborhoods of falsely classified points, *e.g.*, by excluding lines passing through certain regions in 3D space. Similarly, co-occurrence statistics over feature descriptors could be used to filter out irrelevant neighbors.

**Acknowledgements.** This work has been funded by a Google Faculty Research Award, the Chalmers AI Research Centre (Vis-LocLearn), the EU Horizon 2020 project RICAIP (grant agreement No 857306), and the European Regional Development Fund under IMPACT No. CZ.02.1.01/0.0/0.0/15 003/0000468.

## References

- [1] Arandjelović, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: NetVLAD: CNN architecture for weakly supervised place recognition. In: CVPR (2016) [2](#)
- [2] Arth, C., Wagner, D., Klopschitz, M., Irschara, A., Schmalstieg, D.: Wide area localization on mobile phones. In: ISMAR (2009) [1](#)
- [3] Balntas, V., Lenc, K., Vedaldi, A., Mikolajczyk, K.: HPatches: A Benchmark and Evaluation of Handcrafted and Learned Local Descriptors. In: CVPR (2017) [2](#)
- [4] Brachmann, E., Krull, A., Nowozin, S., Shotton, J., Michel, F., Gumhold, S., Rother, C.: DSAC - Differentiable RANSAC for Camera Localization. In: CVPR (2017) [2](#)
- [5] Brachmann, E., Rother, C.: Learning Less is More - 6D Camera Localization via 3D Surface Regression. In: CVPR (2018) [2](#)
- [6] Brachmann, E., Rother, C.: Expert sample consensus applied to camera re-localization. In: ICCV (2019) [2](#)
- [7] Brachmann, E., Rother, C.: Visual camera re-localization from RGB and RGB-D images using DSAC. arXiv:2002.12324 (2020) [2](#)
- [8] Brahmabhatt, S., Gu, J., Kim, K., Hays, J., Kautz, J.: Geometry-aware learning of maps for camera localization. In: CVPR (2018) [2](#)
- [9] Bujnak, M., Kukulova, Z., Pajdla, T.: A general solution to the p4p problem for camera with unknown focal length. In: CVPR (2008) [1](#), [2](#)
- [10] Camposeco, F., Cohen, A., Pollefeys, M., Sattler, T.: Hybrid Scene Compression for Visual Localization. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [1](#)
- [11] Cao, S., Snavely, N.: Minimal Scene Descriptions from Structure from Motion Models. In: CVPR (2014) [1](#)
- [12] Castle, R.O., Klein, G., Murray, D.W.: Video-rate localization in multiple maps for wearable augmented reality. In: ISWC (2008) [1](#)
- [13] Cavallari, T., Bertinetto, L., Mukhoti, J., Torr, P., Golodetz, S.: Let's take this online: Adapting scene coordinate regression network predictions for online rgb-d camera relocalisation. In: 3DV (2019) [2](#)
- [14] Cavallari, T., Golodetz, S., Lord, N.A., Valentin, J., Di Stefano, L., Torr, P.H.S.: On-The-Fly Adaptation of Regression Forests for Online Camera Relocalisation. In: CVPR (2017) [2](#)
- [15] Chen, D.M., Baatz, G., Köser, K., Tsai, S.S., Vedantham, R., Pylvänäinen, T., Roimela, K., Chen, X., Bach, J., Pollefeys, M., Girod, B., Grzeszczuk, R.: City-Scale Landmark Identification on Mobile Devices. In: CVPR (2011) [1](#)
- [16] Choudhary, S., Narayanan, P.J.: Visibility probability structure from sfm datasets and applications. In: ECCV (2012) [1](#), [2](#)
- [17] Chum, O., Matas, J.: Optimal Randomized RANSAC. PAMI **30**(8), 1472–1482 (2008) [1](#)
- [18] Cummins, M., Newman, P.: FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance. IJRR **27**(6), 647–665 (2008). <https://doi.org/10.1177/0278364908090961>, <http://ijr.sagepub.com/cgi/content/abstract/27/6/647> [1](#)
- [19] Donoser, M., Schmalstieg, D.: Discriminative Feature-to-Point Matching in Image-Based Localization. In: CVPR (2014) [2](#)
- [20] Dosovitskiy, A., Brox, T.: Generating images with perceptual similarity metrics based on deep networks. In: Proceedings of the 30th International Conference on Neural Information Processing Systems (2016) [2](#)
- [21] Dosovitskiy, A., Brox, T.: Inverting visual representations with convolutional networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [2](#)
- [22] Dusmanu, M., Rocco, I., Pajdla, T., Pollefeys, M., Sivic, J., Torii, A., Sattler, T.: D2-Net: A trainable CNN for joint detection and description of local features. In: CVPR (2019) [2](#)
- [23] Fischler, M., Bolles, R.: Random Sampling Consensus: A Paradigm for Model Fitting with Application to Image Analysis and Automated Cartography. CACM **24**, 381–395 (1981) [1](#), [2](#)
- [24] Geppert, M., Larsson, V., Speciale, P., Schönberger, J.L., Pollefeys, M.: Privacy Preserving Structure-from-Motion. In: ECCV (2020) [3](#)
- [25] Germain, H., Bourmaud, G., Lepetit, V.: Sparse-To-Dense Hypercolumn Matching for Long-Term Visual Localization. In: International Conference on 3D Vision (3DV) (2019) [2](#)
- [26] Haralick, R., Lee, C.N., Ottenberg, K., Nölle, M.: Review and analysis of solutions of the three point perspective pose estimation problem. IJCV **13**(3), 331–356 (1994) [2](#)
- [27] Hartley, R.I., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge Univ. Press, 2nd edn. (2004) [2](#)
- [28] Heinly, J., Schönberger, J.L., Dunn, E., Frahm, J.M.: Reconstructing the world\* in six days. In: CVPR (2015) [1](#)
- [29] Heng, L., Choi, B., Cui, Z., Geppert, M., Hu, S., Kuan, B., Liu, P., Nguyen, R., Yeo, Y.C., Geiger, A., Lee, G.H., Pollefeys, M., Sattler, T.: Project AutoVision: Localization and 3D Scene Perception for an Autonomous Vehicle with a Multi-Camera System. In: International Conference on Robotics and Automation (ICRA) (2019) [1](#)
- [30] Irschara, A., Zach, C., Frahm, J.M., Bischof, H.: From Structure-from-Motion Point Clouds to Fast Location Recognition. In: CVPR (2009) [1](#), [2](#), [3](#)
- [31] Kahl, F., Hartley, R.: Multiple view geometry under the  $L_\infty$ -norm. PAMI **30**(9), 1603–1617 (2008) [1](#)
- [32] Kato, H., Harada, T.: Image reconstruction from bag-of-visual-words. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014) [2](#)
- [33] Kendall, A., Cipolla, R.: Geometric loss functions for camera pose regression with deep learning. In: CVPR (2017) [2](#)

- [34] Kendall, A., Grimes, M., Cipolla, R.: PoseNet: A Convolutional Network for Real-Time 6-DOF Camera Relocalization. In: ICCV (2015) 2, 6
- [35] Kuiper, N.: Tests concerning random points on a circle. In: Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen. pp. Series A, 63, 38–47 (1960) 5
- [36] Kukulova, Z., Bujnak, M., Pajdla, T.: Real-Time Solution to the Absolute Pose Problem with Unknown Radial Distortion and Focal Length. In: ICCV (2013) 1, 2
- [37] Kukulova, Z., Heller, J., Fitzgibbon, A.: Efficient Intersection of Three Quadrics and Applications in Computer Vision. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) 1, 2
- [38] Larsson, V., Sattler, T., Kukulova, Z., Pollefeys, M.: Revisiting Radial Distortion Absolute Pose. In: The IEEE International Conference on Computer Vision (ICCV) (2019) 1
- [39] Lebeda, K., Matas, J.E.S., Chum, O.: Fixing the Locally Optimized RANSAC. In: British Machine Vision Conference (BMVC) (2012) 1
- [40] Li, Y., Snavely, N., Huttenlocher, D., Fua, P.: Worldwide Pose Estimation Using 3D Point Clouds. In: ECCV (2012) 1, 2, 3
- [41] Li, Y., Snavely, N., Huttenlocher, D.P.: Location Recognition using Prioritized Feature Matching. In: ECCV (2010) 1, 2, 3
- [42] Lim, H., Sinha, S.N., Cohen, M.F., Uyttendaele, M.: Real-Time Image-Based 6-DOF Localization in Large-Scale Environments. In: CVPR (2012) 1
- [43] Lowe, D.: Distinctive Image Features from Scale-Invariant Keypoints. IJCV 60(2) (2004) 1, 2, 3
- [44] Lynen, S., Bosse, M., Furgale, P., Siegwart, R.: Placeless place-recognition. In: International Conference on 3D Vision (3DV). pp. 303–310 (2014) 5
- [45] Lynen, S., Bosse, M., Siegwart, R.: Trajectory based place recognition for efficient large scale localization. In: International Journal for Computer Vision. pp. 49–64 (2017) 5
- [46] Lynen, S., Sattler, T., Bosse, M., Hesch, J., Pollefeys, M., Siegwart, R.: Get Out of My Lab: Large-scale, Real-Time Visual-Inertial Localization. In: RSS (2015) 1
- [47] Lynen, S., Zeisl, B., Aiger, D., Bosse, M., Hesch, J., Pollefeys, M., Siegwart, R., Sattler, T.: Large-scale, real-time visual-inertial localization revisited. Arxiv <https://arxiv.org/abs/1907.00338> (2019), <https://arxiv.org/abs/1907.00338> 1
- [48] Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015) 2
- [49] Massiceti, D., Krull, A., Brachmann, E., Rother, C., Torr, P.H.: Random Forests versus Neural Networks - What's Best for Camera Relocalization? In: ICRA (2017) 2
- [50] Middelberg, S., Sattler, T., Untzelmann, O., Kobbelt, L.: Scalable 6-DOF Localization on Mobile Devices. In: ECCV (2014) 1
- [51] Mur-Artal, R., Tardós, J.D.: ORB-SLAM2: an Open-Source SLAM System for Monocular, Stereo and RGB-D Cameras. TRO 33(5), 1255–1262 (2017) 1
- [52] Nister, D., Stewenius, H.: Scalable recognition with a vocabulary tree. In: CVPR (2006) 5
- [53] Oh, S.J., Benenson, R., Fritz, M., Schiele, B.: Faceless person recognition; privacy implications in social media. In: ECCV. Springer (2016) 2
- [54] Philbin, J., Isard, M., Sivic, J., Zisserman, A.: Descriptor learning for efficient retrieval. In: ECCV (2010) 2
- [55] Pittaluga, F., Koppal, S.J., Kang, S.B., Sinha, S.N.: Revealing Scenes by Inverting Structure From Motion Reconstructions. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 1, 2, 3, 4, 5, 6, 7, 8
- [56] Robertson, D., Cipolla, R.: An image-based system for urban navigation. In: BMVC (2004) 1, 2
- [57] Salas-Moreno, R.F., Newcombe, R.A., Strasdat, H., Kelly, P.H.J., Davison, A.J.: SLAM++: Simultaneous Localisation and Mapping at the Level of Objects. In: CVPR (2013) 1
- [58] Sarlin, P.E., Cadena, C., Siegwart, R., Dymczyk, M.: From Coarse to Fine: Robust Hierarchical Localization at Large Scale. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 1, 2
- [59] Sattler, T., Leibe, B., Kobbelt, L.: Fast Image-Based Localization using Direct 2D-to-3D Matching. In: ICCV (2011) 1
- [60] Sattler, T., Leibe, B., Kobbelt, L.: Efficient & Effective Prioritized Matching for Large-Scale Image-Based Localization. PAMI 39(9), 1744–1756 (2017) 2, 3
- [61] Sattler, T., Havlena, M., Radenovic, F., Schindler, K., Pollefeys, M.: Hyperpoints and fine vocabularies for large-scale location recognition. In: ICCV (2015) 1, 2
- [62] Sattler, T., Torii, A., Sivic, J., Pollefeys, M., Taira, H., Okutomi, M., Pajdla, T.: Are Large-Scale 3D Models Really Necessary for Accurate Visual Localization? In: CVPR (2017) 1, 2
- [63] Sattler, T., Weyand, T., Leibe, B., Kobbelt, L.: Image Retrieval for Image-Based Localization Revisited. In: BMVC (2012) 1
- [64] Sattler, T., Zhou, Q., Pollefeys, M., Leal-Taixe, L.: Understanding the limitations of cnn-based absolute camera pose regression. In: CVPR (2019) 2
- [65] Schönberger, J.L., Frahm, J.M.: Structure-From-Motion Revisited. In: CVPR (June 2016) 1
- [66] Schwesinger, U., Buerki, M., Timpner, J., Rottman, S., Wolf, L., Paz, L.M., Grimmet, H., Posner, I., Newman, P., Häne, C., Heng, L., Lee, G.H., Sattler, T., Pollefeys, M., Allodi, M., Valenti, F., Mimura, K., Goebelsmann, B., Siegwart, R.: Automated Valet Parking and Charging for e-Mobility: Results of the V-Charge Project. In: IV (2016) 1
- [67] Shibuya, M., Sumikura, S., Sakurada, K.: Privacy Preserving Visual SLAM. In: ECCV (2020) 1, 3

- [68] Shotton, J., Glocker, B., Zach, C., Izadi, S., Criminisi, A., Fitzgibbon, A.: Scene Coordinate Regression Forests for Camera Relocalization in RGB-D Images. In: CVPR (2013) [2](#)
- [69] Snavely, N., Seitz, S., Szeliski, R.: Modeling the World from Internet Photo Collections. IJCV **80**(2), 189–210 (2008) [1](#)
- [70] Song, Z., Chen, W., Campbell, D., Li, H.: Deep Novel View Synthesis from Colored 3D Point Clouds. In: ECCV (2020) [1, 2, 3](#)
- [71] Speciale, P., Schonberger, J.L., Kang, S.B., Sinha, S.N., Pollefeys, M.: Privacy preserving image-based localization. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [1, 2, 3, 5, 6, 7, 8](#)
- [72] Speciale, P., Schonberger, J.L., Sinha, S.N., Pollefeys, M.: Privacy Preserving Image Queries for Camera Localization. In: The IEEE International Conference on Computer Vision (ICCV) (2019) [3](#)
- [73] Strecha, C., von Hansen, W., Van Gool, L., Fua, P., Thoennessen, U.: On benchmarking camera calibration and multi-view stereo for high resolution imagery. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2008) [6](#)
- [74] Svärm, L., Enqvist, O., Kahl, F., Oskarsson, M.: City-Scale Localization for Cameras with Known Vertical Direction. PAMI **39**(7), 1455–1461 (2017) [1, 2](#)
- [75] Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A.: InLoc: Indoor visual localization with dense matching and view synthesis. In: CVPR (2018) [2](#)
- [76] Toft, C., Maddern, W., Torii, A., Hammarstrand, L., Stenborg, E., Safari, D., Okutomi, M., Pollefeys, M., Sivic, J., Pajdla, T., Kahl, F., Sattler, T.: Long-term visual localization revisited. PAMI (2020) [1](#)
- [77] Torii, A., Arandjelović, R., Sivic, J., Okutomi, M., Pajdla, T.: 24/7 Place Recognition by View Synthesis. In: CVPR (2015) [2](#)
- [78] Torii, A., Sivic, J., Pajdla, T.: Visual localization by linear combination of image descriptors. In: Proceedings of the 2nd IEEE Workshop on Mobile Vision, with ICCV (2011) [2](#)
- [79] Valentin, J., Dai, A., Niessner, M., Kohli, P., Torr, P., Izadi, S., Keskin, C.: Learning to Navigate the Energy Landscape. In: International Conference on 3D Vision (3DV) (2016) [6](#)
- [80] Vondrick, C., Khosla, A., Malisiewicz, T., Torralba, A.: Hoggles: Visualizing object detection features. In: IEEE International Conference on Computer Vision (CVPR) (2013) [2](#)
- [81] Walch, F., Hazirbas, C., Leal-Taixé, L., Sattler, T., Hilsenbeck, S., Cremers, D.: Image-Based Localization Using LSTMs for Structured Feature Correlation. In: ICCV (2017) [2](#)
- [82] Weinzaepfel, P., Jégou, H., Pérez, P.: Reconstructing an image from its local descriptors. In: CVPR (2011) [2](#)
- [83] Wu, Z., Wang, Z., Wang, Z., Jin, H.: Towards privacy-preserving visual recognition via adversarial training: A pilot study. In: ECCV. Springer (2018) [2](#)
- [84] Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., Lipson, H.: Understanding Neural Networks Through Deep Visualization. In: Deep Learning Workshop, International Conference on Machine Learning (ICML) (2015) [2](#)
- [85] Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: ECCV (2014) [2](#)
- [86] Zeisl, B., Sattler, T., Pollefeys, M.: Camera pose voting for large-scale image-based localization. In: ICCV (2015) [1, 2, 3](#)
- [87] Zhang, W., Kosecka, J.: Image based localization in urban environments. In: 3DPVT (2006) [1, 2](#)
- [88] Zhou, Q., Sattler, T., Pollefeys, M., Leal-Taixé, L.: To learn or not to learn: Visual localization from essential matrices. arXiv:1908.01293 (2019) [2](#)