

Robust Point Cloud Based Reconstruction of Large-Scale Outdoor Scenes

Ziquan Lan Zi Jian Yew Gim Hee Lee
 Department of Computer Science, National University of Singapore
 {ziquan, zijian.yew, gimhee.lee}@comp.nus.edu.sg

Abstract

Outlier feature matches and loop-closures that survived front-end data association can lead to catastrophic failures in the back-end optimization of large-scale point cloud based 3D reconstruction. To alleviate this problem, we propose a probabilistic approach for robust back-end optimization in the presence of outliers. More specifically, we model the problem as a Bayesian network and solve it using the Expectation-Maximization algorithm. Our approach leverages on a long-tail Cauchy distribution to suppress outlier feature matches in the odometry constraints, and a Cauchy-Uniform mixture model with a set of binary latent variables to simultaneously suppress outlier loop-closure constraints and outlier feature matches in the inlier loop-closure constraints. Furthermore, we show that by using a Gaussian-Uniform mixture model, our approach degenerates to the formulation of a state-of-the-art approach for robust indoor reconstruction. Experimental results demonstrate that our approach has comparable performance with the state-of-the-art on a benchmark indoor dataset, and outperforms it on a large-scale outdoor dataset. Our source code can be found on the project website¹.

1. Introduction

Point cloud reconstruction of outdoor scenes has many important applications such as 3D architectural modeling, terrestrial surveying, Simultaneous Localization and Mapping (SLAM) for autonomous vehicles, *etc.* Compared to images, point clouds from 3D scanners exhibit less variation under different weather or lighting conditions, *e.g.*, summer and winter (Fig. 1), or day and night (Fig. 5). Furthermore, the depths of point clouds from 3D scanners are more accurate than image-based reconstructions. Consequently, point clouds from 3D scanners are preferred for large-scale outdoor 3D reconstructions. Most existing methods for 3D reconstruction are solved via a two-step approach: a front-end data association step and a back-end optimization step. More specifically, data association is used to establish fea-

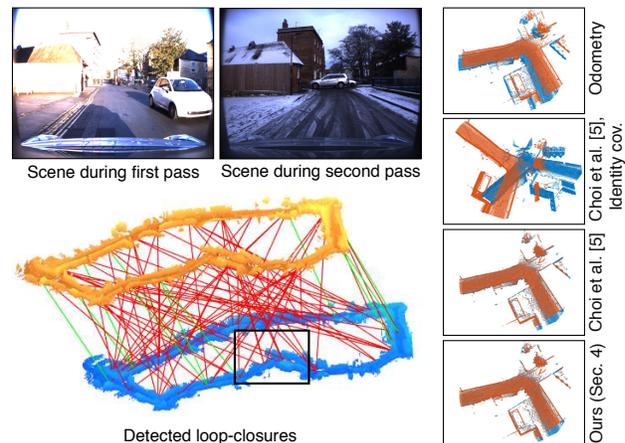


Figure 1. Reconstruction of a 1km route traversed in two different seasons: summer (orange) and winter (blue). The outlier (red links) loop-closures significantly outnumber the inliers (green links). Four zoomed-in point clouds on the right are reconstructed from different methods.

ture matches [30] in point cloud fragments for registration, and loop-closures [26] between point cloud fragments for pose-graph [21] optimization. Unfortunately, no existing algorithm for feature matching and loop-closure detection guarantees complete elimination of outliers. Although outlier feature matches are usually handled with RANSAC-based geometric verification [16, 30], such pairwise checks do not consider global consistency. In addition, the numerous efforts on improving the accuracy in loop-closure detection [6, 8, 20, 26] are not completely free from false positives. Many back-end optimization algorithms [13, 14, 21] are based on non-linear least-squares that lack the robustness to cope with outliers. A small number of outliers would consequently lead to catastrophic failures in the 3D reconstructions. Several prior works focus on disabling outlier loop-closures in the back-end optimization [5, 15, 25]. However, these methods do not consider the effect from the outlier feature matches with the exception of [34] that solves global geometric registration in a very small-scale problem setting.

The main contribution of this paper is a probabilistic

¹<https://github.com/ziquan111/RobustPCLReconstruction>

approach for robust back-end optimization to handle outliers from a weak front-end data association in large-scale point cloud based reconstructions. Our approach simultaneously suppresses outlier feature matches and loop-closures. To this end, we model our robust point cloud reconstruction problem as a Bayesian network. The global poses of the point cloud fragments are the unknown parameters, and odometry and loop-closure constraints are the observed variables. A binary latent variable is assigned to each loop-closure constraint; it determines whether a loop-closure constraint is an inlier or outlier. We model feature matches in the odometry constraints with a long-tail Cauchy distribution to gain robustness to outlier matches. Additionally, we use a Cauchy-Uniform mixture model for loop-closure constraints. The uniform and Cauchy distributions model outlier loop-closures and the feature matches in inlier loop-closures, respectively. In contrast to many existing back-end optimizers that use rigid transformations as the odometry and loop-closure constraints [5, 14, 15, 21, 25], we use the distances between feature matches to exert direct influence on these matches.

We use the Expectation-Maximization (EM) algorithm [3, 15] to find the globally consistent poses of the point cloud fragments (Sec. 4). The EM algorithm iterates between the Expectation and Maximization steps. In the Expectation step, the posterior of a loop-closure constraint being an inlier is updated. In the Maximization step, a local optimal solution for the global poses is found from maximizing the expected complete data log-likelihood over the posterior from the expectation step. We also generalize our approach to solve reconstruction problems with an easier setting (Sec. 5). In particular, a strong assumption is imposed: odometry and inlier loop-closure constraints are free from outlier feature matches. We show that by using a Gaussian-Uniform mixture model, our approach degenerates to the formulation of a state-of-the-art approach for robust indoor reconstruction [5]. Fig. 1 shows an example of the reconstruction result with our method compared to other methods in the presence of outliers.

2. Related Work

Reconstruction of outdoor scenes has been studied in [22, 23]. Schöps et al. [23] propose a set of filtering steps to detect and discard unreliable depth measurements acquired from a RGB-D camera. However, loop-closures is not detected and this can lead to reconstruction failures. Relying on very accurate GPS/INS, Pollefeys et al. [22] propose a 3D reconstruction system from RGB images. However, GPS/INS signal may be unavailable or unreliable, especially on cloudy days or in urban canyons. Our work relies on neither GPS/INS nor RGB images. In contrast, we focus on reconstruction from point cloud data acquired from 3D scanners that is less sensitive to weather or light-

ing changes. There are also many works on indoor scene reconstruction. Since the seminal KinectFusion [18], there are several follow-up algorithms [4, 19, 27]. Unfortunately, these methods do not detect loop-closures. Nonetheless, there are many RGB-D reconstruction methods with loop-closure detection [5, 7, 10, 11, 24, 28, 31, 32, 33].

Choi et al. [5] achieve the state-of-the-art performance for indoor reconstruction with robust loop-closure. However, they assume no outlier feature matches in the odometry and inlier loop-closure constraints. We relax this assumption to achieve robust feature matching. More specifically, [5] estimates a switch variable [25] for each loop-closure constraint using line processes [2]. Outlier loop-closures are disabled by setting the respective switch variables to zero. Additional switch prior terms are imposed and chosen empirically [25] to prevent a trivial solution of removing all loop-closure constraints. In comparison, our approach does not require the additional prior terms. We estimate the posterior of a loop-closure being an inlier constraint in the Expectation step shown in Sec. 4. The EM approach is also used by Lee et al. [15]. However, they solve a robust pose-graph optimization problem without coping with the feature matches for reconstruction.

3. Overview

In this section, we provide an overview of our reconstruction pipeline that consists of four main components: point cloud fragment construction, point cloud registration, loop-closure detection, and robust reconstruction with EM.

Point cloud fragment construction. A single scan from a 3D scanner, *e.g.* LiDAR, contains limited number of points. We integrate multiple consecutive scans with odometry readings obtained from dead reckoning *e.g.*, the Inertial Navigation System (INS) [17] to form local point cloud fragments. A set of 3D features is then extracted from each point cloud fragment using [30].

Point cloud registration. The top k_1 feature matches between two consecutive point cloud fragments F_i and F_{i+1} are retained as the odometry constraint $X_{i,i+1}$. Since consecutive fragments overlap sufficiently by construction [17, 30], we define $X_{i,i+1}$ as a reliable constraint but note that it can contain outlier feature matches.

Loop-closure detection. It is inefficient to perform an exhaustive pairwise registration for large-scale outdoor scenes with many point cloud fragments. Hence, we perform point cloud based place-recognition [26] to identify a set of candidate loop-closures. We retain the top k_2 potential loop-closures for each fragment and remove the duplicates. For each loop-closure between fragments F_i and F_j , we keep

the set of top k_1 feature matches denoted as Y_{ij} . We define Y_{ij} as a loop-closure constraint, which can either be an inlier or outlier. Similar to the odometry constraint, an inlier loop-closure Y_{ij} can also contain outlier feature matches.

Robust reconstruction with EM. The constraints from point cloud registration and loop-closure detection can contain outliers. In particular, both odometry and loop-closure constraints can contain outlier feature matches. Moreover, many detected loop-closures are false positives. In the next section, we describe our probabilistic modeling approach to simultaneously suppress outlier feature matches and false loop-closures. The EM algorithm is used to solve for the globally consistent fragment poses. Optional refinement using ICP can be applied to further improve the global point cloud registration.

4. Robust Reconstruction with EM

We model the robust reconstruction problem as a Bayesian network shown in Fig. 2. Let $T = [T_1, \dots, T_i, \dots, T_N]^\top$, where $T_i \in \text{SE}(3)$, denote the N fragment poses, $X = [X_{12}, \dots, X_{i,i+1}, \dots, X_{N-1,N}]^\top$ denote the $N-1$ odometry constraints obtained in point cloud registration, and $Y = [\dots, Y_{ij}, \dots]^\top$ denote the M loop-closure constraints obtained in loop-closure detection. We explicitly assign the loop-closure constraints into 2 clusters that represent the inliers and outliers. For each loop-closure constraint Y_{ij} , we introduce a corresponding assignment variable $Z_{ij} = [Z_{ij,in}, Z_{ij,out}]^\top \in \{[1, 0]^\top, [0, 1]^\top\}$. Z_{ij} is a one-hot vector: $Z_{ij,in} = 1$ and $Z_{ij,out} = 1$ assigns Y_{ij} as an inlier and outlier loop-closure constraint, respectively. We use $Z = [\dots, Z_{ij}, \dots]^\top$ to denote the assignment variables. T is the unknown parameter, Z is the latent variable, and X and Y are both observed variables.

Robust reconstruction can be solved as finding the Maximum a Posterior (MAP) solution of $p(T|X, Y)$. However, the MAP solution involves an intractable step of marginalization over the latent variable Z . We circumvent this problem by using the EM algorithm that takes the maximization of the expected complete data log-likelihood over the posterior of the latent variables. The EM algorithm iterates between the Expectation and Maximization steps. In the Expectation step, we use T^{old} , *i.e.*, fragment poses solved from the previous iteration to find the posterior distribution of the latent variable Z ,

$$p(Z|Y, T^{old}) = \frac{p(Y|Z, T^{old})p(Z|T^{old})}{p(Y|T^{old})}, \quad (1)$$

in which Z does not depend on X , since they are conditionally independent given Y according to the Bayesian network in Fig. 2.

In the Maximization step, the posterior distribution (Eq. (1)) is used to update T by maximizing the expectation of

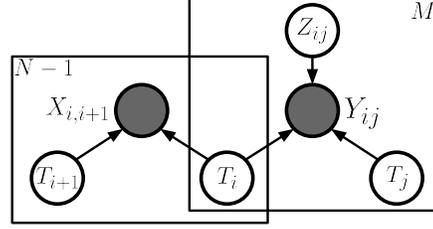


Figure 2. Bayesian network representation of the robust reconstruction problem. T_{i+1} , T_i and T_j are fragment poses. $X_{i,i+1}$ is an odometry constraint. Y_{ij} is a loop-closure constraint. Z_{ij} is an assignment variable. $N-1$ and M indicate the numbers of odometry constraints and loop-closure constraints respectively.

the complete data log-likelihood denoted by

$$\begin{aligned} Q_{EM} &:= \sum_Z p(Z|Y, T^{old}) \ln p(X, Y, Z|T) \\ &= \underbrace{\ln p(X|T)}_{Q_X} + \underbrace{\sum_Z p(Z|Y, T^{old}) \ln p(Y, Z|T)}_{Q_Y}. \end{aligned} \quad (2)$$

We define Q_X for the term with odometry constraints, and Q_Y for the term with loop-closure constraints.

Initialization. The unknown parameters, *i.e.*, global poses T of the N fragments, are initialized with the relative poses computed from odometry constraints X using ICP. Other dead reckoning methods such as wheel odometry and/or INS readings can also be used.

4.1. Modeling Odometry Constraints

Odometry constraints are obtained from point cloud registration between two consecutive point cloud fragments. Recall that an odometry constraint $X_{i,i+1}$ is a set of feature matches between fragments F_i and F_{i+1} , which can contain outlier matches. To gain robustness, we model each feature match $(\mathbf{p}, \mathbf{q}) \in X_{i,i+1}$ with a long-tail multivariate Cauchy distribution. Suppose these feature matches are independent and identically distributed (*i.i.d.*), we take a geometric mean over their product to get

$$p(X_{i,i+1}|T) = \left(\prod_{(\mathbf{p}, \mathbf{q}) \in X_{i,i+1}} \text{Cauchy}_{i,i+1}(\mathbf{p}, \mathbf{q}) \right)^{\frac{1}{|X_{i,i+1}|}}, \quad (3)$$

where

$$\text{Cauchy}_{i,i+1}(\mathbf{p}, \mathbf{q}) = \frac{1}{\pi^2 \sqrt{\det \Sigma} (1 + d_{\Sigma}^2(T_i \mathbf{p}, T_{i+1} \mathbf{q}))^2}, \quad (4)$$

which we assume an isotropic covariance $\Sigma = \sigma^2 \mathbf{I}$ with scale σ , and d_{Σ} denotes the Mahalanobis distance such that

$$d_{\Sigma}^2(T_i \mathbf{p}, T_{i+1} \mathbf{q}) = (T_i \mathbf{p} - T_{i+1} \mathbf{q})^\top \Sigma^{-1} (T_i \mathbf{p} - T_{i+1} \mathbf{q}). \quad (5)$$

The value of σ is set based on the density of extracted features. For example, $\sigma = 0.5\text{m}$ in the outdoor dataset.

4.2. Modeling Loop-Closure Constraints

A loop-closure constraint Y_{ij} is the set of feature matches between fragments F_i and F_j . We propose to use a Cauchy-Uniform mixture model to cope with the (1) outlier loop-closure constraints and (2) outlier feature matches in the inlier loop-closure constraints.

To distinguish between inlier and outlier loop-closures, we model the distribution of assignment variable Z as a Bernoulli distribution defined by the inlier probability $\lambda \in [0, 1]$,

$$p(Z_{ij}) = \lambda^{Z_{ij, \text{in}}} (1 - \lambda)^{Z_{ij, \text{out}}}. \quad (6)$$

Next, we use two distributions: Cauchy and Uniform distributions to model the inlier and outlier loop-closure constraints, respectively.

Cauchy distribution – inlier loop-closure constraints. The inlier loop-closure constraints can contain outlier feature matches. We use the same multivariate Cauchy distribution as Eq. (3) and further reorganize the terms. We define C_{ij} for brevity, such that

$$C_{ij} := p(Y_{ij} | Z_{ij, \text{in}} = 1, T) = \pi^{-2} \sigma^{-3} e^{-2A_{ij}}, \quad (7)$$

in which

$$A_{ij} = \frac{1}{|Y_{ij}|} \sum_{(\mathbf{p}, \mathbf{q}) \in Y_{ij}} \ln\left(1 + \frac{\|T_i \mathbf{p} - T_j \mathbf{q}\|^2}{\sigma^2}\right), \quad (8)$$

and $|Y_{ij}|$ denotes the number of feature matches in Y_{ij} .

Uniform distribution – outlier loop-closure constraints.

We model the outlier loop-closure constraints with a uniform distribution defined by a constant probability $u \in (0, 1)$,

$$p(Y_{ij} | Z_{ij, \text{out}} = 1, T) = u. \quad (9)$$

4.3. Expectation Step

Recall that the expectation step is evaluated in Eq. (1). Plugging Eq. (6), (7) and (9) into the Bayes' formula, we obtain the posterior of being an inlier loop-closure constraint,

$$\mathcal{P}_{ij}^{\text{in}} := p(Z_{ij, \text{in}} = 1 | Y, T) = \frac{\Theta}{\Theta + e^{2A_{ij}}}, \quad (10)$$

where

$$\Theta = \frac{\lambda}{(1 - \lambda)u\pi^2\sigma^3}. \quad (11)$$

The constant Θ consists of two distribution parameters: λ is the probability of being an inlier loop-closure; u is the

constant probability to uniformly sample a random loop-closure, which are difficult to set manually based on different datasets. Hence, we propose to estimate Θ based on the input data. More specifically, we learn Θ from the odometry constraints, since all odometry constraints are effectively inlier loop-closure constraints.

The process to learn Θ is as follows. First, for each odometry constraint $X_{i, i+1}$, we denote its corresponding error term $m_{i, i+1} = e^{2A_{i, i+1}}$ (analogous to Eq. (10)), where

$$A_{i, i+1} = \frac{1}{|X_{i, i+1}|} \sum_{(\mathbf{p}, \mathbf{q}) \in X_{i, i+1}} \ln\left(1 + \frac{\|T_i \mathbf{p} - T_{i+1} \mathbf{q}\|^2}{\sigma^2}\right). \quad (12)$$

Next, we compute the median error denoted as \hat{m} . Since we regard all odometry constraints as inlier loop-closure constraints, let

$$\frac{\Theta}{\Theta + \hat{m}} = \hat{p}, \quad (13)$$

where we set $\hat{p} = 90\%$, meaning that a loop-closure Y_{ij} with a small error ($e^{2A_{ij}} < \hat{m}$) is very likely to be an inlier ($\mathcal{P}_{ij}^{\text{in}} > \hat{p}$). Finally, we solve for Θ using Eq. (13).

4.4. Maximization Step

In the maximization step, we solve for T that maximizes $Q_{EM} = Q_X + Q_Y$, where Q_X and Q_Y are shorthand notations defined in Eq. (2). These two terms are evaluated independently, and then optimized jointly.

Evaluate Q_X . Assuming the odometry constraints in X are i.i.d., the joint probability of all odometry constraints is given by

$$p(X|T) = \prod_{i=1}^{N-1} p(X_{i, i+1}|T). \quad (14)$$

Substituting the joint probability of the feature matches within each odometry constraint (Eq. (3)), we can rewrite Q_X as

$$Q_X = -2 \sum_{i=1}^{N-1} A_{i, i+1} + \text{const}. \quad (15)$$

Evaluate Q_Y . Using the product rule, the joint probability of loop-closure constraints and their corresponding assignment variables can be written as $p(Y, Z|T) = p(Z)p(Y|Z, T)$. Plugging Eq. (6), (7) and (9) in, we have

$$p(Y, Z|T) = \prod_{i, j} (\lambda C_{ij})^{Z_{ij, \text{in}}} ((1 - \lambda)u)^{Z_{ij, \text{out}}}. \quad (16)$$

We can rewrite Q_Y as

$$Q_Y = \sum_{i, j} \mathcal{P}_{ij}^{\text{in}} \ln C_{ij} + \text{const}, \quad (17)$$

with the joint probability from Eq. (16) and the posterior from Eq. (10), which can be further expanded to

$$Q_Y = -2 \sum_{i,j} \mathcal{P}_{ij}^{in} A_{ij} + \text{const.} \quad (18)$$

Maximize Q_{EM} . The maximization of $Q_X + Q_Y$ can be reformulated into a non-linear least-squares problem with the following objective function

$$\begin{aligned} \operatorname{argmin}_T \sum_{i,j} \frac{\mathcal{P}_{ij}^{in}}{|Y_{ij}|} \sum_{(\mathbf{p}, \mathbf{q}) \in Y_{ij}} \ln\left(1 + \frac{\|T_i \mathbf{p} - T_j \mathbf{q}\|^2}{\sigma^2}\right) \quad (19) \\ + \sum_{i=1}^{N-1} \frac{1}{|X_{i,i+1}|} \sum_{(\mathbf{p}, \mathbf{q}) \in X_{i,i+1}} \ln\left(1 + \frac{\|T_i \mathbf{p} - T_{i+1} \mathbf{q}\|^2}{\sigma^2}\right), \end{aligned}$$

which can be easily optimized using the sparse Cholesky solver in Google Ceres [1]. The computation complexity is cubic to the total number of feature matches.

5. Generalization using EM

In the previous section, we solved the problem when constraints are contaminated with outlier feature matches. In this section, we study a problem with an easier setting where correct loop-closure constraints contain no outlier feature matches. Recall that long-tail multivariate Cauchy distribution is used to gain robustness against outlier feature matches. We replace the multivariate Cauchy distribution with a multivariate Gaussian distribution for the easier problem without outlier feature matches, and show that our EM formulation degenerates to the formulation of a state-of-the-art approach for robust indoor reconstruction [5]. To avoid repetition, we only highlight the major differences to the previous section. Each analogous term is augmented with a superscript G that stands for ‘‘Gaussian’’.

Odometry constraints. Replacing the multivariate Cauchy distribution in Eq. (3) with a multivariate Gaussian distribution, we have

$$p^G(X_{i,i+1}|T) = \left(\prod_{(\mathbf{p}, \mathbf{q}) \in X_{i,i+1}} \text{Gauss}_{i,i+1}(\mathbf{p}, \mathbf{q}) \right)^{\frac{1}{|X_{i,i+1}|}}, \quad (20)$$

where

$$\text{Gauss}_{i,i+1}(\mathbf{p}, \mathbf{q}) = \frac{\exp\left(-\frac{1}{2} d_{\Sigma}^2(T_i \mathbf{p}, T_{i+1} \mathbf{q})\right)}{\sqrt{(2\pi)^3 \det \Sigma}}, \quad (21)$$

and Σ and d_{Σ} remain unchanged.

Loop-closure constraints. We note that the Bernoulli distribution in Eq. (6) still holds, and the major changes

start from Eq. (7). Using the multivariate Gaussian distribution, we have

$$G_{ij} := p^G(Y_{ij}|Z_{ij, in} = 1, T) = (\sqrt{2\pi}\sigma)^{-3} e^{-\frac{B_{ij}}{2\sigma^2}}, \quad (22)$$

in which

$$B_{ij} = \frac{1}{|Y_{ij}|} \sum_{(\mathbf{p}, \mathbf{q}) \in Y_{ij}} \|T_i \mathbf{p} - T_j \mathbf{q}\|^2, \quad (23)$$

and $|Y_{ij}|$ is the number of feature matches. We note that B_{ij} is a sum-of-square errors that can lead to arithmetic overflow in the $e^{B_{ij}}$ term from the posterior of the latent variable Z_{ij} (analogous to Eq. (10)). In contrast, there is no arithmetic overflow in the $e^{A_{ij}}$ term from Eq. (10) since A_{ij} from Eq. (8) is a sum-of-log errors. We propose to alleviate the arithmetic overflow problem by using a Pareto distribution that approximates G_{ij} as

$$G_{ij} \approx \frac{x_0}{B_{ij}^2}, \quad (24)$$

where $x_0 > 0$ is a scale parameter. For outlier loop-closures, the uniform distribution in Eq. (9) still holds.

Expectation step. Using the approximation of G_{ij} in Eq. (24), the posterior from Eq. (10) becomes

$$\mathcal{P}_{ij}^{inG} \approx \frac{\Theta^G}{\Theta^G + B_{ij}^2}, \quad (25)$$

where

$$\Theta^G = \frac{x_0 \lambda}{(1 - \lambda)u}. \quad (26)$$

It becomes apparent in \mathcal{P}_{ij}^{inG} that the arithmetic overflow problem is alleviated by the replacement of $e^{B_{ij}}$ with B_{ij}^2 . In the previous section, Θ in Eq. (13) is learned from the median error \hat{m} of all the error terms $m_{i,i+1} = e^{2A_{i,i+1}}$ in the odometry constraints. Unfortunately, the median error \hat{m}^G from $m_{i,i+1}^G = B_{i,i+1}^2$ becomes uninformative because we assume no outlier feature matches, *i.e.*, $\hat{m}^G \rightarrow 0$ since $m_{i,i+1}^G = B_{i,i+1}^2 \rightarrow 0$. Despite the absence of outlier feature matches, $\|T_i \mathbf{p} - T_j \mathbf{q}\|$ is upper bounded by some threshold, ϵ . Hence, the mean error term can be directly estimated from Eq. (23) as $\hat{m}^G = \epsilon^2$. Subsequently, let

$$\frac{\Theta^G}{\Theta^G + \hat{m}^G} = \hat{p}, \quad (27)$$

where we set $\hat{p} = 90\%$ and solve for Θ^G . We set $\epsilon = 0.05\text{m}$ for our experiments on the indoor dataset (see next section) based on the typical magnitude of sensor noise.

		Living room 1	Living room 2	Office 1	Office 2	Average
Before pruning	Recall(%)	61.2	49.7	64.4	61.5	59.2
	Precision(%)	27.2	17.0	19.2	14.9	19.6
Choi et al. [5] after pruning	Recall(%)	57.6	49.7	63.3	60.7	57.8
	Precision(%)	95.1	97.4	98.3	100.0	97.7
Ours (Sec. 5) after pruning	Recall(%)	58.7	48.4	63.9	61.5	58.1
	Precision(%)	97.0	94.9	96.6	93.6	95.4

Table 1. Results of robust optimization on the indoor dataset. Our method shows comparable result with the state-of-the-art.

	Living room 1	Living room 2	Office 1	Office 2	Average
Whelan et al. [27]	0.22	0.14	0.13	0.13	0.16
Kerl et al. [12]	0.21	0.06	0.11	0.10	0.12
SUN3D [29]	0.09	0.07	0.13	0.09	0.10
Choi et al. [5]	0.04	0.07	0.03	0.04	0.05
Ours (Sec. 5)	0.06	0.09	0.05	0.04	0.06
GT Trajectory	0.04	0.04	0.03	0.03	0.04

Table 2. Reconstruction accuracy on the indoor dataset. The entries are the mean distances of each model to its respective ground-truth surface (in meters). Our proposed method shows comparable result with the state-of-the-art and outperforms the rest.

Maximization step. Finally, we reformulate the maximization problem as a non-linear least-squares problem with the following objective function

$$\begin{aligned} \operatorname{argmin}_T \sum_{i,j} \frac{\mathcal{P}_{ij}^{inG}}{|Y_{ij}|} \sum_{(\mathbf{p},\mathbf{q}) \in Y_{ij}} \|T_i \mathbf{p} - T_j \mathbf{q}\|^2 \quad (28) \\ + \sum_{i=1}^{N-1} \frac{1}{|X_{i,i+1}|} \sum_{(\mathbf{p},\mathbf{q}) \in X_{i,i+1}} \|T_i \mathbf{p} - T_{i+1} \mathbf{q}\|^2, \end{aligned}$$

which is similar to the formulation in [5] with two minor differences. First, we average the square errors over the number of feature matches but [5] does not. Second, we estimate the posterior \mathcal{P}_{ij}^{inG} by iterating between the Expectation and Maximization steps but [5] estimates it using line processes [2]. It is important to note that Eq. (28) is derived from the original Gaussian formulation in Eq. (22) instead of the Pareto approximation in Eq. (24).

6. Evaluation

We use the experimental results from two datasets for the comparison between our approach and the state-of-the-art approach [5]. The first dataset is from small-scale indoor scenes with no outlier feature matches in the odometry and inlier loop-closure constraints, and the second dataset is from large-scale outdoor scenes with outlier feature matches. Our Gaussian-Uniform EM (Sec. 5) and Cauchy-Uniform EM (Sec. 4) are evaluated on the small-scale indoor and large-scale outdoor datasets, respectively.

6.1. Small-Scale Indoor Scenes

The ‘‘Augmented ICL-NUIM Dataset’’ provided and augmented by [9] and [5], respectively, is used as the small-

scale indoor dataset. This dataset is generated from synthetic indoor environments and includes two models: a living room and an office. There are two RGB-D image sequences for each model, resulting in a total of four test cases. To ensure fair comparison, we follow the same evaluation criteria and experimental settings as [5].

Results. Tab. 1 shows the comparison of the average recall and precision of the loop-closures on (1) before pruning, (2) [5] after pruning and (3) our method after pruning. Here, ‘‘before pruning’’ refers the loop-closures from the loop-closure detection, and ‘‘after pruning’’ refers to the inlier loop-closures after robust optimization. It can be seen that the average precision and recall of our method is comparable to [5]. This is an expected result since we showed in Sec. 5 that our method degenerates to the method in [5] with minor differences in the absence of outlier feature matches. We further evaluate the reconstruction accuracy of the final model using the error metric proposed in [9], *i.e.*, the mean distance of the reconstructed surfaces to the ground truth surfaces. Tab. 2 shows the comparison of the reconstruction accuracy of our method to other existing approaches. In addition, as suggested in [5], the reconstruction accuracy of the model obtained from fusing the input depth images with the ground truth trajectory (denoted as GT Trajectory in Tab. 2) is reported for reference. As expected, our method shows comparable result with the state-of-the-art on the indoor dataset.

6.2. Large-Scale Outdoor Scenes

The large-scale outdoor dataset is based on the ‘‘Oxford Robotcar Dataset’’ [17]. It consists of 3D point clouds captured with a LiDAR sensor mounted on a car that repeatedly

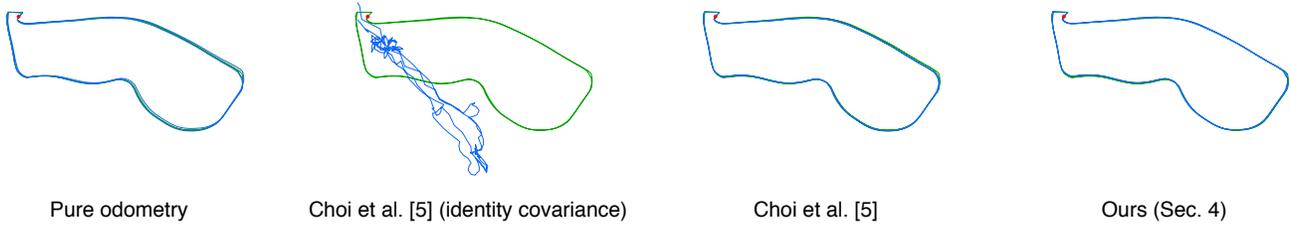


Figure 3. Trajectories on the 1km route. Each trajectory (blue) is overlaid with the GPS/INS trajectory (green). Red asterisk indicates the position of the 1st fragment pose.

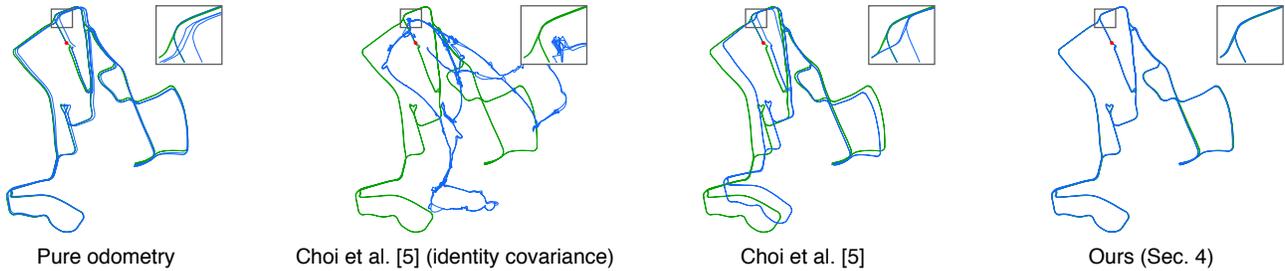


Figure 4. Trajectories on the city-scale route. Each trajectory (blue) is overlaid with the GPS/INS trajectory (green). A zoomed-in region is shown on the top right corner for each trajectory. Red asterisk indicates the position of the 1st fragment pose.

drives through Oxford, UK, at different times over a year. We select two different driving routes from the dataset, a short route (about 1km) and a long route (city-scale). Furthermore, we take two traversals at different times for each route, resulting four traversals in total. Unlike the synthetic indoor dataset, there is no ground truth of the surface geometry. We evaluate the trajectory accuracy against the GPS/INS readings as an indirect measurement of reconstruction accuracy. We prepare the dataset as follows:

- **Point cloud fragments.** We integrate the push-broom 2D LiDAR scans and their corresponding INS readings into the 3D point clouds. We segment the data into fragments with 30m radius for every 10m interval. Each fragment is then downsampled using a VoxelGrid filter with a grid size of 0.2m. 242 and 1770 fragments are constructed for the 1km route and the city-scale route, respectively.
- **Odometry trajectory.** The odometry trajectory is disconnected due to discontinuous INS data since we are combining two traversals. We simulate the odometry trajectory via geometric registrations between consecutive point cloud fragments, and manually identify one linkage transformation between the two traversals. We also check the entire odometry trajectory to ensure that there are no remaining erroneous transformations. The resulting odometry trajectory is used to initialize the fragment poses, T .
- **Odometry constraints.** For every two consecutive frames along the odometry trajectory, we perform point cloud registration as described in Sec. 3. Specifically, we

extract 1024 features for each fragment, and collect the top 200 feature matches to form an odometry constraint. Note that the feature matches are selected without additional geometric verification, and it can contain outliers. 241 and 1769 odometry constraints are constructed for the 1km route and the city-scale route, respectively.

- **Loop-closure constraints.** We perform loop-closure detection as described in Sec. 3. We take every 5th fragment along the trajectory as a keyframe fragment; loop-closures are detected among the selected keyframe fragments. For the 1km route, we find the top 5 loop-closures for each keyframe fragment and then remove the duplicates. For the city-scale route, we find the top 10 loop-closures for each keyframe fragment and then remove the duplicates. 171 and 1438 loop-closure constraints are constructed for the 1km and city-scale route, respectively. The outlier loop-closure ratio is more than 80% for both routes.

Baseline Methods. We compare the effectiveness of our approach with two baseline methods based on [5]: a stronger and a weaker baseline. The stronger baseline encodes uncertainty information of the feature matches between two fragments into a covariance matrix. The feature matches used to construct the covariance matrix are those within 1m apart after geometric registration. Refer to [5] for the more details on the covariance matrix. The covariance matrix of the weaker baseline is set to identity, *i.e.*, no uncertainty information on the feature matches. The relative poses between the point cloud fragments computed from

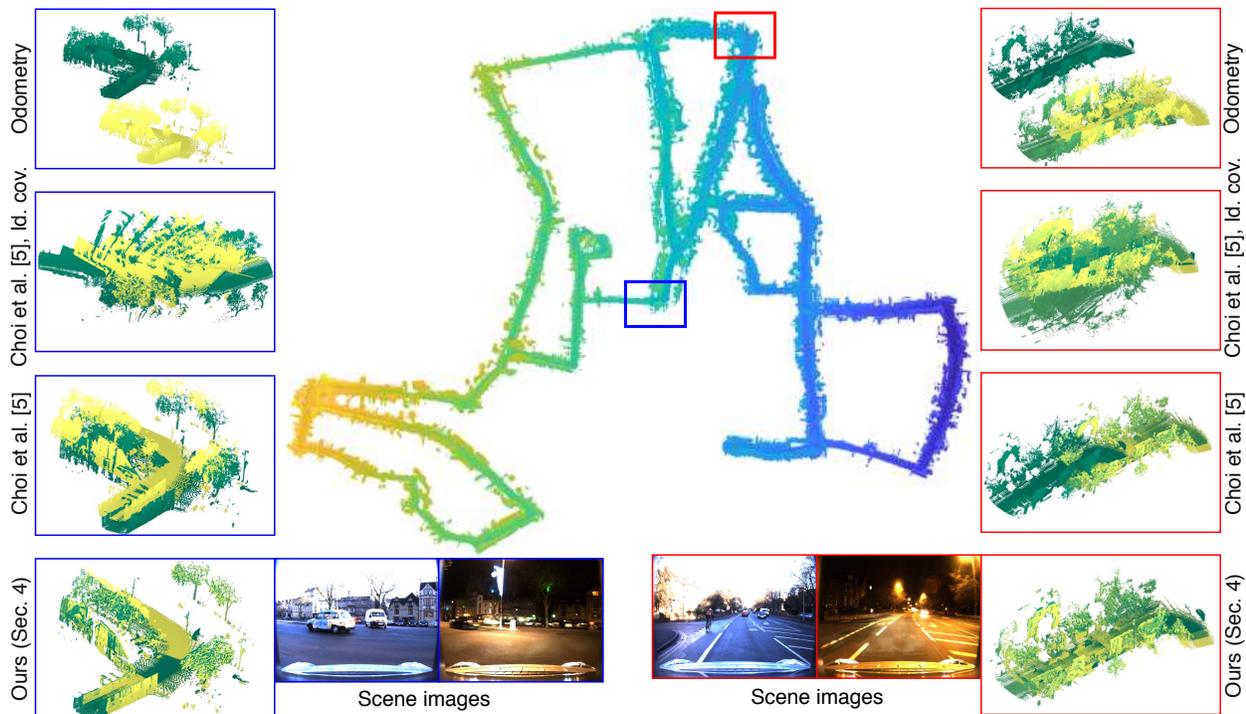


Figure 5. Reconstruction results on the city-scale route traversed in day and night. The two columns of zoomed-in point clouds are reconstructed based on different trajectories.

	1km	City-scale
Odometry	1.85	11.81
Choi et al. [5] (identity covariance)	123.24	207.93
Choi et al. [5]	1.97	50.92
Ours (Sec. 4)	1.34	2.45

Table 3. Reconstruction accuracy on outdoor dataset. Each entry is the mean distance of the estimated poses to the GPS/INS ground truth (in meters).

ICP are used as the odometry and loop-closure constraints in the baseline methods.

Results. Tab. 3 summarizes the mean distances of the estimated poses to the GPS/INS trajectory as an indirect measure of the reconstruction accuracy on the 1km and city-scale outdoor datasets. Fig. 3 and 4 show the plots of the trajectories. We align the first five fragment poses with the GPS/INS trajectory, error measurements start after the 5th fragment pose. The results show that the accuracy increases when more information about the feature matches is considered in the optimization process. We can see from Tab. 3, and Fig. 3 and 4 that the weaker baseline

([5] with uninformative identity covariance) without information of the feature matches gives the worst performance. The stronger baseline ([5] with informative covariance matrix) that encodes information of feature matches using the covariance matrix shows better performance. In contrast, our method that directly takes feature matches as the odometry and loop-closure constraints outperforms the two baselines. Furthermore, Fig. 1 and 5 show reconstruction results for qualitative evaluation. It can be seen from the bottom left and right plots in Fig. 5 that our method produces the sharpest reconstructions of the 3D point clouds.

7. Conclusion

In this paper, we proposed a probabilistic approach for robust point cloud reconstruction of large-scale outdoor scenes. Our approach leverages on a Cauchy-Uniform mixture model to simultaneously suppress outlier feature matches and loop-closures. Moreover, we showed that by using a Gaussian-Uniform mixture model, our approach degenerates to the formulation of a state-of-the-art approach for robust indoor reconstruction. We verified our proposed methods on both indoor and outdoor benchmark datasets.

Acknowledgement

This work is supported in part by a Singapore MOE Tier 1 grant R-252-000-A65-114.

References

- [1] S. Agarwal, K. Mierle, and Others. Ceres solver. <http://ceres-solver.org>.
- [2] M. J. Black and A. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. In *IJCV*, 1996.
- [3] G. Celeux and G. Govaert. A classification em algorithm for clustering and two stochastic versions. In *CSDA*, 1992.
- [4] J. Chen, D. Bautembach, and S. Izadi. Scalable real-time volumetric surface reconstruction. In *TOG*, 2013.
- [5] S. Choi, Q.-Y. Zhou, and V. Koltun. Robust reconstruction of indoor scenes. In *ICCV*, 2015.
- [6] M. Cummins and P. Newman. Appearance-only slam at large scale with fab-map 2.0. In *IJRR*, 2011.
- [7] F. Endres, J. Hess, J. Sturm, D. Cremers, and W. Burgard. 3-d mapping with an rgb-d camera. In *T-RO*, 2014.
- [8] D. Galvez-Lopez and J. D. Tardos. Real-time loop detection with bags of binary words. In *IROS*, 2011.
- [9] A. Handa, T. Whelan, J. McDonald, and A. J. Davison. A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM. In *ICRA*, 2014.
- [10] P. Henry, D. Fox, A. Bhowmik, and R. Mongia. Patch volumes: Segmentation-based consistent mapping with rgb-d cameras. In *3DV*, 2013.
- [11] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox. Rgb-d mapping: Using kinect-style depth cameras for dense 3d modeling of indoor environments. In *IJRR*, 2012.
- [12] C. Kerl, J. Sturm, and D. Cremers. Dense visual slam for rgb-d cameras. In *IROS*, 2013.
- [13] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger. Factor graphs and the sum-product algorithm. In *T-IT*, 2001.
- [14] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g 2 o: A general framework for graph optimization. In *ICRA*, 2011.
- [15] G. H. Lee, F. Fraundorfer, and M. Pollefeys. Robust pose-graph loop-closures with expectation-maximization. In *IROS*, 2013.
- [16] G. H. Lee and M. Pollefeys. Unsupervised learning of threshold for geometric verification in visual-based loop-closure. In *ICRA*, 2014.
- [17] W. Maddern, G. Pascoe, C. Linegar, and P. Newman. 1 year, 1000 km: The oxford robotcar dataset. In *IJRR*, 2017.
- [18] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *ISMAR*, 2011.
- [19] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger. Real-time 3d reconstruction at scale using voxel hashing. In *TOG*, 2013.
- [20] D. Nister and H. Stewenius. Scalable recognition with a vocabulary tree. In *CVPR*, 2006.
- [21] E. Olson, J. Leonard, and S. Teller. Fast iterative alignment of pose graphs with poor initial estimates. In *ICRA*, 2006.
- [22] M. Pollefeys, D. Nistér, J.-M. Frahm, A. Akbarzadeh, P. Mordohai, B. Clipp, C. Engels, D. Gallup, S.-J. Kim, P. Merrell, et al. Detailed real-time urban 3d reconstruction from video. In *IJCV*, 2008.
- [23] T. Schöps, T. Sattler, C. Häne, and M. Pollefeys. Large-scale outdoor 3d reconstruction on a mobile device. In *CVIU*, 2017.
- [24] F. Steinbrucker, C. Kerl, and D. Cremers. Large-scale multi-resolution surface reconstruction from rgb-d sequences. In *ICCV*, 2013.
- [25] N. Sünderhauf and P. Protzel. Switchable constraints for robust pose graph slam. In *IROS*, 2012.
- [26] M. A. Uy and G. H. Lee. PointNetVLAD: Deep Point Cloud Based Retrieval for Large-Scale Place Recognition. In *CVPR*, 2018.
- [27] T. Whelan, H. Johannsson, M. Kaess, J. J. Leonard, and J. McDonald. Robust real-time visual odometry for dense rgb-d mapping. In *ICRA*, 2013.
- [28] T. Whelan, M. Kaess, J. J. Leonard, and J. McDonald. Deformation-based loop closure for large scale dense rgb-d slam. In *IROS*, 2013.
- [29] J. Xiao, A. Owens, and A. Torralba. Sun3d: A database of big spaces reconstructed using sfm and object labels. In *ICCV*, 2013.
- [30] Z. J. Yew and G. H. Lee. 3DFeat-Net: Weakly Supervised Local 3D Features for Point Cloud Registration. In *ECCV*, 2018.
- [31] Q.-Y. Zhou and V. Koltun. Dense scene reconstruction with points of interest. In *TOG*, 2013.
- [32] Q.-Y. Zhou and V. Koltun. Simultaneous localization and calibration: Self-calibration of consumer depth cameras. In *CVPR*, 2014.
- [33] Q.-Y. Zhou, S. Miller, and V. Koltun. Elastic fragments for dense scene reconstruction. In *ICCV*, 2013.
- [34] Q.-Y. Zhou, J. Park, and V. Koltun. Fast global registration. In *ECCV*, 2016.