2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

CBGL: Fast Monte Carlo Passive Global Localisation of 2D LIDAR Sensor

Alexandros Filotheou

***Abstract*— Navigation of a mobile robot is conditioned on the knowledge of its pose. In observer-based localisation configurations its initial pose may not be knowable in advance, leading to the need of its estimation. Solutions to the problem of global localisation are either robust against noise and environment arbitrariness but require motion and time, which may (need to) be economised on, or require minimal estimation time but assume environmental structure, may be sensitive to noise, and demand preprocessing and tuning. This article proposes a method that retains the strengths and avoids the weaknesses of the two approaches. The method leverages properties of the Cumulative Absolute Error per Ray (CAER) metric with respect to the errors of pose hypotheses of a 2D LIDAR sensor, and utilises scan–to–map-scan matching for fine(r) pose estimations. A large number of tests, in real and simulated conditions, involving disparate environments and sensor properties, illustrate that the proposed method outperforms state-of-the-art methods of both classes of solutions in terms of pose discovery rate and execution time. The source code is available for download.**

***Index Terms*— global localisation, 2D LIDAR, monte carlo, scan–to–map-scan matching**

I. Introduction

The knowledge of a mobile robot's initial pose is a prerequisite in tasks involving its navigation, especially in contexts where an observer is used for pose tracking. However, the robot's pose is not predictable or pre-settable in all conditions. This lack of predictability necessitates ad hoc estimation without prior information. Various sensor and map modalities have been investigated in the literature: from 2D and 3D LIDAR sensors [1], [2], RGBD cameras [3], and RFID equipment [4], to keyframe-based submaps [5] and metric maps [6]. In practice the latter combined with 2D LIDAR sensors have become the de facto means of mobile robot localisation and navigation due to the sensor's high measurement precision and frequency, almost no need for preprocessing, and low cost compared to 3D LIDAR sensors.

This article addresses the problem of Passive Global Localisation of a 2D LIDAR sensor in a 2D metric map, i.e. the estimation of its location and orientation within the map, under complete locational and orientational uncertainty, without prescribing robot motion commands for further knowledge acquisition. The problem is formalised in Problem P:

Problem P. Let the unknown pose of an immobile 2D range sensor whose angular range is λ be $\boldsymbol{p}(\boldsymbol{l}, \theta)$, $\boldsymbol{l} = (x, y)$, with respect to the reference frame of map $\boldsymbol{M}$. Let the range sensor measure range scan $\mathcal{S}_R$. The objective is the estimation of $\boldsymbol{p}$ given $\mathcal{S}_R$, $\boldsymbol{M}$, and λ.

The author is with the Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, 54124 Thessaloniki, Greece alefilot@auth.gr

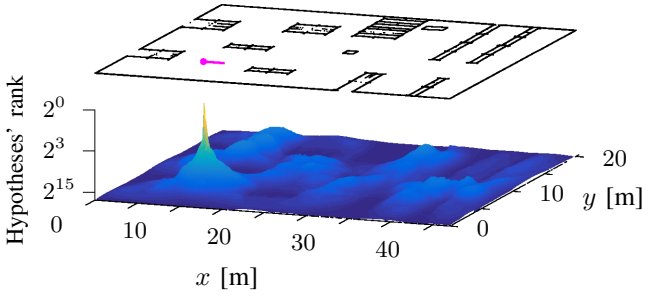

Fig. 1: Top: the map of an environment and the pose of a panoramic 2D LIDAR sensor (magenta). Given a LIDAR sensor's 2D measurement, at its core, CBGL disperses pose hypotheses within the map and ranks them ascendingly (bottom) according to the value of the Cumulative Absolute Error per Ray metric (Eq. (1)). This ranking may estimate the pose of the sensor (a) quickly due to the metric's low computational complexity, and (b) accurately due to (i) proportionality between the pose estimate error and the value of the metric for pose estimates in a neighbourhood of the sensor's pose, and (ii) lack of disproportionality outside of that neighbourhood (Fig. 2)

For the solution to Problem P this article introduces CBGL: a single-shot Monte Carlo method (a) which makes no assumptions regarding the structure or the particulars of the sensor's environment or the sensor's characteristics, (b) whose pose errors exhibit robustness to varying sensor angular range, and (c) which operates with three optionally-set and intuitive parameters, which trade accuracy for execution time. The central contributions of the article are:

- To the best of the author's knowledge the fastest Monte Carlo global localisation method that employs a 2D LIDAR that achieves higher pose discovery rates than state-of-the-art methods
- The extension and validation of the Cumulative Absolute Error per Ray metric's ability to estimate pose error hierarchies solely from real and virtual range scans, extended from scan–to–map-scan matching (`sm2`) during pose-tracking, where pose estimates are "few" ($\leq 2^5$) and their errors are "small", to the context of global localisation, where the set of hypotheses and their errors may be arbitrarily large
- The thorough evaluation of the proposed method against (a) established state-of-the-art localisation methods, (b) real and publicly available benchmark conditions, and (c) varying characteristics of environments and sensors, which target real conditions, that pose hindrances to global localisation methods

The rest of the article is structured as follows: Section II provides necessary definitions and the notation employed. Section III gives a brief review of solutions to problem P, and the relation of the proposed method to them. The latter's methodology is presented in section IV, its evaluation in section V, and its limitations in section VI. Section VII concludes this study.

II. Definitions and Notation

Let $\mathcal{A} = \{\alpha_i : \alpha_i \in \mathbb{R}\}$, $i \in \mathtt{I} = \langle 0, 1, \ldots, n-1\rangle$, denote a set of n elements, $\langle\cdot\rangle$ denote an ordered set, $\mathcal{A}_{\uparrow}$ the set $\mathcal{A}$ ordered in ascending order, the bracket notation $\mathcal{A}[\mathtt{I}] = \mathcal{A}$ denote indexing, and notation $\mathcal{A}_{k:l}$, $0 \leq k \leq l$, denote limited indexing: $\mathcal{A}_{k:l} = \{\mathcal{A}[k], \mathcal{A}[k+1], \ldots, \mathcal{A}[l]\}$.

Definition I. *Pose and pose estimate*.—The object of localisation is the estimation of the—fundamentally unknown—3DoF pose $\boldsymbol{p}$ of a sensor or robot: $\hat{\boldsymbol{p}} = (\hat{\boldsymbol{l}}, \hat{\theta})$, where $\hat{\boldsymbol{l}} = (\hat{x}, \hat{y}) \in \mathbb{R}^2$ is its location on the 2D plane and $\hat{\theta} \in [-\pi, \pi)$ rad its orientation relative to the positive x axis.

Definition II. *Range scan captured from a 2D LIDAR sensor*.—A range scan $\mathcal{S}$, measured by a 2D LIDAR sensor, consisting of N_s rays over an angular range $\lambda \in (0, 2\pi]$, is an ordered map of angles to distances between objects and the sensor within its radial range $r_{\max}$: $\mathcal{S} : \Theta \to \mathbb{R}_{\geq 0}$, $\Theta = \{\theta_n \in [-\frac{\lambda}{2}, +\frac{\lambda}{2}) : \theta_n = -\frac{\lambda}{2} + \lambda\frac{n}{N_s},\ n = 0, 1, \ldots, N_s - 1\}$. Angles θ_n are expressed relative to the sensor's heading, in the sensor's frame of reference [7].

Definition III. *Map-scan*.—A map-scan is a virtual scan that encapsulates the same pieces of information as a scan derived from a physical sensor (Def. II). Contrary to the latter it refers to distances within a map $\boldsymbol{M}$ of an environment rather than within the environment itself; the latter is represented by $\boldsymbol{M}$ usually in Occupancy Grid or ordered Point Cloud form. A map-scan $\mathcal{S}_V^{\boldsymbol{M}}(\hat{\boldsymbol{p}})$ is derived by locating intersections of rays emanating from a sensor's pose estimate $\hat{\boldsymbol{p}}$ and the boundaries of $\boldsymbol{M}$.

Definition IV. *The Cumulative Absolute Error per Ray metric*.—Let $\mathcal{S}_p$ and $\mathcal{S}_q$ be two range scans (Defs. II, III), equal in angular range λ and size N_s. The value of the Cumulative Absolute Error per Ray (CAER) metric $\psi \in \mathbb{R}_{\geq 0}$ between $\mathcal{S}_p$ and $\mathcal{S}_q$ is given by

$$\psi(\mathcal{S}_p, \mathcal{S}_q) \triangleq \sum_{n=0}^{N_s-1} \Big|\mathcal{S}_p[\theta_n] - \mathcal{S}_q[\theta_n]\Big| \tag{1}$$

Definition V. *Pose densities*.—The locational and angular density, $d_{\boldsymbol{l}}, d_{\alpha} \in \mathbb{N}$, of a pose set $\mathcal{H} = \{\hat{\boldsymbol{p}}_i\}$ respectively denote the mean number of pose estimates dispersed on a given map per unit area of space and per angular cycle.

Definition VI. *Admissibility of solution*.—A pose estimate $\hat{\boldsymbol{p}}(\hat{\boldsymbol{l}}, \hat{\theta}) \in \mathbb{R}^2 \times [-\pi, +\pi)$, may be deemed an admissible solution to Problem P iff $\|\boldsymbol{p} - \hat{\boldsymbol{p}}\|_2 \leq \delta_0$ and $\|\boldsymbol{l} - \hat{\boldsymbol{l}}\|_2 \leq \delta_{\boldsymbol{l}}$ and $|\theta - \hat{\theta}| \leq \delta_\theta$, where $\delta_{\boldsymbol{l}}, \delta_\theta, \delta_0 \in \mathbb{R}_{>0}$: $\delta_{\boldsymbol{l}}^2 + \delta_\theta^2 = \delta_0^2$.

III. Related Work

The literature concerning the solution to the problem of global localisation with the use of a 2D LIDAR sensor is rich. A recent and comprehensive survey on global localisation may be found in [8], while a brief overview is given below.

In broad terms global localisation approaches may be divided into two categories: (a) approaches that operate in feature space, that is, methods that extract features from measurements and the map and establish correspondences between them, and (b) approaches that directly exploit only raw measurements. In the latter category a number of methods solve the problem in an iterative Bayesian Monte Carlo fashion, i.e. by dispersing hypotheses within the map and updating the belief of the robot's pose by incorporating new measurements as it moves until estimate convergence [9]–[13]. Although motion may be assistive—or even required—in certain conditions, such as those of repetitive map structures, the requirement of motion (a) may give rise to safety concerns (the robot may not even be visible), and (b) increases estimation time. The Monte Carlo method proposed in this article operates directly in measurement space as well but, in contrast, is not iterative, and does not require motion or more than one measurement. As a result CBGL is a single-shot global localisation approach that is able to process more hypotheses in less time, resulting in an improvement in the number of correctly estimated locations, and making it able to compete against (traditionally faster) feature-based approaches in terms of execution time.

Research on the feature-based approaches has been more extensive due to the richness, adaptability, and efficacy of methods originated in the computer vision field, and their low execution time. Relevant methods mainly perform detection of key-points in a measurement, followed by the calculation of a distinctive signature, which is then matched to similarly- and pre-computed place-signatures [1], [14]–[20]. In principle, however, unstructured environments cannot be relied upon for the existence of features due to their absence or their sparse and fortuitous distribution (although Deep Neural Network approaches have demonstrated increased performance in place recognition with the use of 3D LIDARs [21]–[23]). Structured environments, on the other hand, manifest different features depending on the particularities of the environment, where features may be present but not in a sufficiently undisturbed state due to sensor noise or map-to-environment mismatch. Furthermore, feature-based methods require the tuning of parameters in a per-environment basis, which hinders the range and degree of their applicability and efficacy. In contrast CBGL (a) makes no assumptions on- and exploits no environmental structure, and (b) uses three parameters, whose setting is optional and intuitive.

In sum the proposed method retains most of the positive qualities of the two main approaches to global localisation and avoids their pitfalls: it uses multiple hypotheses for robustness against uncertainty, assumes no motion, environmental structure, or parameter tuning in a per-environment or sensor basis, and is robust against noise. The motivation

of CBGL originates from seeking to achieve a greater degree of universality, reliability, and portability across multiple and disparate environments, by aiming to economise on the use of resources (environmental assumptions; number, type, and cost of sensors; number of measurements; time). The proposed method is most akin to the two tested methods in [24]: all three are single-shot 2D LIDAR-based Monte Carlo approaches, but CBGL computes a measure of the alignment of the measurement to the virtual scan captured from each hypothesis *before* scan–to–map-scan matching, which occurs subsequently and only for a small subset of the virtual scans most aligned to the measurement. This, along with the low computational complexity of the alignment measure that CBGL utilises (Def. IV), greatly diminish its execution time. So far this alignment measure has only been tested against the context of `sm2` in pose tracking, i.e. for pose hypotheses in a neighbourhood of the sensor's true pose [25], [26].

IV. Methodology

A. Motivation

The proposed method's motivation lies in the simple but powerful facts that are in general exhibited through fig. 2: Let a pose estimate of a 2D LIDAR sensor be within the given map (Def. I); then the value of the CAER metric (Def. IV) between the scan measured by the sensor (Def. II) and the map-scan captured from the estimate (Def. III) is (a) proportional to both the estimate's location error and orientation error in a neighbourhood of the sensor's true pose, (b) not inversely proportional to them outside of it, and (c) greater in value outside than inside of it. In other words comparing any two estimates of the sensor's pose in terms of the CAER metric, where at least one is in a neighbourhood of it, is enough to establish a pose error hierarchy between them. This is significant because: while the error of a single pose estimate is unknowable, the top of a hierarchy of dense-enough pose estimates (Def. V) ordered by CAER values may provide the neighbourhood of the sensor's pose, and hence an admissible estimate of the pose itself (Def. VI). Moreover, the CAER metric is calculable from the assumptions of Problem P, with low computational complexity $\mathcal{O}(N_s)$. The relationships of proportionality between the CAER metric of a pose estimate and its location and orientation errors have been discovered and successfully exploited in the context of non-global localisation, and specifically in pose-tracking, for the production of lidar odometry [25] and the reduction of localisation's pose estimate error [26]. In that context estimate errors are close to the origin, in contrast to the distribution of hypotheses' errors in global localisation methods. Fig. 11 (top) shows the relations between the total errors of hypotheses and their (a) CAER values (left) and (b) ranks when ordered by CAER ascending (right), which resulted from an experimental procedure similar to that which produced fig. 2. The above motivate investigation on whether the CAER metric could prove equally beneficial in settings more uncertain than pose-tracking, that is, where the location and orientation of pose estimates extend farther away from the sensor's true location and orientation.

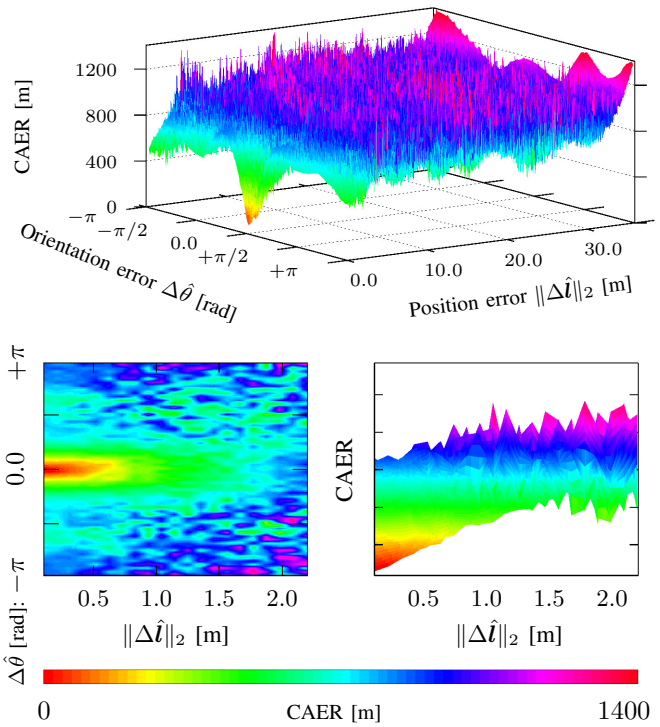

Fig. 2: Top: without loss of generality, a typical plot of the Cumulative Absolute Error per Ray metric (Eq. (1)) of 10^6 hypotheses disperesed in the map of environment WAREHOUSE (Fig. 1 and section V-B). Bottom: focused view on hypotheses with location error close to the origin

B. The CBGL Algorithm

In order to test the efficacy of the CAER metric in aiding the solution of Problem P, we introduce the CAER-Based Global Localisation algorithm (CBGL), described in block diagram form in fig. 4 and in pseudocode in Algorithm I.[1] Given map $\boldsymbol{M}$, CBGL first generates a set of pose hypotheses $\mathcal{H}$ (Def. I). Their positions are randomly generated uniformly within the map's traversable space; their orientations within $[-\pi, \pi)$ rad. Then from these poses CBGL computes map-scans (Def. III). Given these map-scans and a LIDAR's 2D measurement $\mathcal{S}_R$ (Def. II), CBGL subsequently computes the CAER value (Def. IV) associated with each pose hypothesis. It then ranks them in ascending order and selects the k estimates with the least CAER values in an attempt to estimate the k hypotheses with the least pose error, producing set $\mathcal{H}_1$ (Alg. II). Estimation in this sense rests on the motivation of subsection IV-A. Fig. 1 (bottom) indicatively depicts such a pose hierarchy.

One challenge is choosing such k, $d_{\boldsymbol{l}}$, and d_α (Def. V) that, given pose estimate error requirements $\delta_{\boldsymbol{l}}$, δ_θ, CBGL produces admissible pose estimates (Def. VI) while being executed in timely manner. Given the method's Monte Carlo nature, optimistically, the only option for increasing the accuracy of the final pose estimate by a factor of two would be to double locational density $d_{\boldsymbol{l}}$. Instead of doing that—and thereby doubling the method's execution time—

[1]For a rigorous mathematical formulation of the hypothesis underpinning CBGL see [27].

subsequent to the calculation of set $\mathcal{H}_1$, CBGL scan–to–map-scan matches the map-scans captured from the pose estimates of $\mathcal{H}_1$ against the range scan measured by the real sensor (sm2) [26], [28], producing pose set $\mathcal{H}_2$ (Alg. III).

Matching allows for (a) the correction of the pose of true positive estimates, (b) by the same token the potential divergence of spurious, false positive pose estimates, and hence their elimination as pose estimate candidates, (c) the decoupling of the final pose estimate's error from the method's chosen densities, and (d) the production of finer pose estimates without excessive increase in execution time. The principle of sm2 is depicted in fig. 3.

CBGL's output is the pose estimate with the least CAER value within the group of k matched estimates.

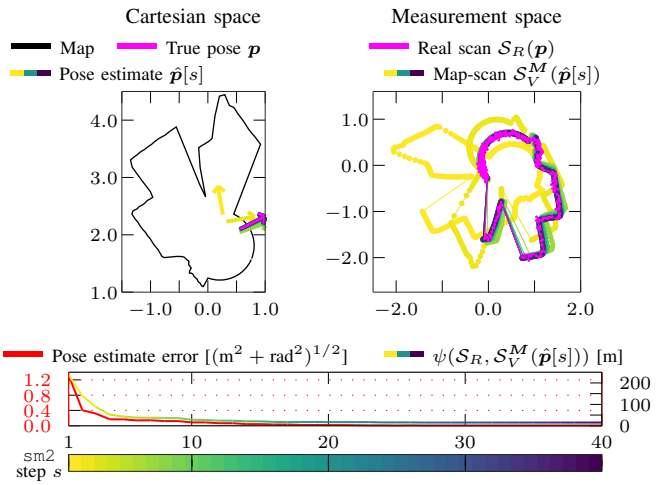

Fig. 3: Depiction of the process of scan–to–map-scan matching (sm2) in principle, which takes place for all $\hat{\boldsymbol{p}}_i = \mathcal{H}_1[i]$ (Alg. I, line 11), $i = 0, 1, \ldots, k-1$. The registration of map-scan $\mathcal{S}_V^{\boldsymbol{M}}(\hat{\boldsymbol{p}}[s])$ to $\mathcal{S}_R$ at step s results in pose $\hat{\boldsymbol{p}}[s+1]$, where $\hat{\boldsymbol{p}}_i[0] = \mathcal{H}_1[i]$. The error of $\hat{\boldsymbol{p}}[s+1]$ is, in principle, reduced compared to $\hat{\boldsymbol{p}}[s]$, and the next map-scan is captured from it. The process is iterative and completes upon estimate convergence

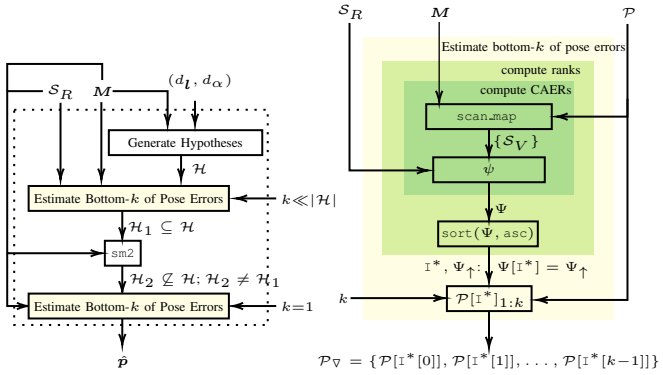

Fig. 4: CBGL in block diagram form. Left: Given map $\boldsymbol{M}$, CBGL first generates a set of pose hypotheses $\mathcal{H}$. Then it estimates the k hypotheses with the least pose error (right; Alg. II). As a final step, it scan–to–map-scan matches these to $\mathcal{S}_R$ for finer estimation (Alg. III). CBGL's output pose estimate is that with the minimum CAER among the k matched estimates. See Algs. I, II, III for notation

Algorithm I CBGL

Input: $\mathcal{S}_R$, λ, $\boldsymbol{M}$, $(d_{\boldsymbol{l}}, d_\alpha)$, k
Output: Pose estimate of sensor measuring range scan $\mathcal{S}_R$

1: $A \leftarrow$ calculate_area(free($\boldsymbol{M}$))
2: $\mathcal{H}, \mathcal{H}_1, \mathcal{H}_2 \leftarrow \{\varnothing\}$
3: **for** $i \leftarrow 0, 1, \ldots, d_{\boldsymbol{l}} \cdot A - 1$ **do**
4: $\quad(\hat{x}, \hat{y}, \hat{\theta}) \leftarrow$ rand(): $(x, y) \in$ free($\boldsymbol{M}$), $\hat{\theta} \in [-\pi, +\pi)$
5: $\quad$**for** $j \leftarrow 0, 1, \ldots, d_\alpha - 1$ **do**
6: $\quad\quad\mathcal{H} \leftarrow \{\mathcal{H}, (\hat{x}, \hat{y}, \hat{\theta} + j \cdot 2\pi/d_\alpha)\}$
7: $\quad$**end for**
8: **end for**
9: $\mathcal{H}_1 \leftarrow$ bottom_k_poses($\mathcal{S}_R, \boldsymbol{M}, \mathcal{H}, k, \lambda$) (Alg. II)
10: **for** $i \leftarrow 0, 1, \ldots, |\mathcal{H}_1| - 1$ **do**
11: $\quad\mathcal{H}_2 \leftarrow \{\mathcal{H}_2,$ sm2$(\mathcal{S}_R, \lambda, \boldsymbol{M}, \mathcal{H}_1[i])\}$ (Alg. III—e.g. x1 [26])
12: **end for**
13: **return** bottom_k_poses($\mathcal{S}_R, \boldsymbol{M}, \mathcal{H}_2, 1, \lambda$)

Algorithm II bottom_k_poses

Input: $\mathcal{S}_R$, $\boldsymbol{M}$, $\mathcal{P}$, k, λ
Output: Set of k poses of $\mathcal{P}$ with least CAER values, $\mathcal{P}_\triangledown$

1: $\Psi \leftarrow \{\varnothing\}$
2: **for** $q \leftarrow 0, 1, \ldots, |\mathcal{P}| - 1$ **do**
3: $\quad\mathcal{S}_V^q \leftarrow \mathcal{S}_V^{\boldsymbol{M}}(\mathcal{P}[q]) =$ scan_map($\boldsymbol{M}, \mathcal{P}[q], \lambda$)
4: $\quad\Psi \leftarrow \{\Psi,$ CAER$(\mathcal{S}_R, \mathcal{S}_V^q)\}$ (Eq. (1))
5: **end for**
6: $[\Psi_\uparrow, \mathtt{I}^*] \leftarrow$ sort(Ψ, asc), such that $\Psi[\mathtt{I}^*] = \Psi_\uparrow$
7: **return** $\mathcal{P}_\triangledown = \{\mathcal{P}[\mathtt{I}^*[0]], \mathcal{P}[\mathtt{I}^*[1]], \ldots, \mathcal{P}[\mathtt{I}^*[k-1]]\}$

Algorithm III sm2

Input: $\mathcal{S}_R$, λ, $\boldsymbol{M}$, $\hat{\boldsymbol{p}}$
Output: $\hat{\boldsymbol{p}}$ + correction that aligns $\mathcal{S}_V^{\boldsymbol{M}}(\hat{\boldsymbol{p}})$ to $\mathcal{S}_R$

1: $\mathcal{S}_V \leftarrow$ scan_map($\boldsymbol{M}, \hat{\boldsymbol{p}}, \lambda$)
2: $\boldsymbol{\Delta p} \leftarrow$ scan–match($\mathcal{S}_R, \mathcal{S}_V$) (e.g. ICP [29], FSM [25])
3: **return** $\hat{\boldsymbol{p}} + \boldsymbol{\Delta p}$

V. Experimental Evaluation

This section focuses on evaluating the performance of state of the art methods in global localization and CBGL in addressing Problem P in static environments, varying environmental conditions, sensor configurations, and map representations. CBGL's parameters are set to $(d_{\boldsymbol{l}}, d_\alpha, k) = (40, 2^5, 10)$ after initial tests with the real dataset used in subsection V-A. The rationale of choosing appropriate $d_{\boldsymbol{l}}$ and d_α is depicted in fig. 5; k is chosen as such in order to retain a high-enough true positive discovery rate without significant increase in execution time. The locational threshold $\delta_{\boldsymbol{l}} = 0.5$ m is used as a tighter solution admissibility criterion than that in [24]; the thresold itself was determined through experimental procedure with a YDLIDAR TG30 sensor (footnote 3; ibid). References to sets $\mathcal{H}_*$ are made to fig. 4 and lines 6, 9, and 11 of Alg. I. Tests are performed with a processor of 12 threads and clock speed 4.00 GHz.

A. Experiments in real conditions

The first type of test is conducted using a Hokuyo UTM-30LX sensor, whose angular range $\lambda = 3\pi/2$ rad and radial range $r_{\max} = 30.0$ m, in the Electrical and Computer Engineering Department's Laboratory of Computer Systems Architecture (CSAL), of the Aristotle University of Thessaloniki, an Occupancy Grid map of which is depicted in fig. 6. The map's resolution is 0.01 m / pixel, and it was constructed via ROS package `open_karto`. The sensor was mounted on a Robotnik RB1 robot, which was teleoperated within the environment while scans were being recorded. This resulted in $N_E = 6669$ range scans, whose number of rays are downsampled by a factor of four before being inputted to CBGL and Advanced Localization System (ALS) [30]. The latter implements Free-Space Features [1] and, contrary to CBGL, is not a single-shot method; however, it is selected for comparison against CBGL due to the fact that it is the only state of the art global localisation method which exhibits feasible execution times with respect to the collected range dataset's volume (see bottom of fig. 8a for indicative execution times of state of the art methods). CBGL's internal `sm2` method is chosen to be PLICP [31] due to its low execution time and the sensor's non-panoramic field of view. The top of fig. 7 depicts the proportion of output pose estimates from each method whose position and orientation error is lower than outlier thresholds δ_l, δ_θ; at the bottom they are depicted exclusively for CBGL's outputs and its internal pose sets. Table I provides a summary of the errors and execution times of the two methods.

From the experimental evidence it is clear that (a) CBGL manages to produce admissible solutions to Problem P 991 times out of a thousand for an outliers' locational threshold $\delta_l = 0.5$ m when an angular threshold δ_θ is not considered, and (b) CBGL outperforms ALS in terms of (i) position and orientation errors, (ii) number of pose estimates within all locational and angular thresholds, and (iii) execution time.

	Position Err. [m]		Orientation Err. [rad]		Exec. Time [sec]	
	Mean	std	Mean	std	Mean	std
ALS	0.500	0.265	1.956	1.167	6.15	5.32
CBGL	0.041	0.045	0.011	0.019	1.61	0.06

TABLE I: Mean and standard deviation of (a) errors of ALS and CBGL with regard to position and orientation and (b) their execution times, with regard to experiments in real conditions

B. Simulations against sources of uncertainty

The second type of test concerns the main limiting factor of global localisation methods, i.e. uncertainty—: arising e.g. from spurious measurements, repeatability of surroundings, missing or corrupted range information, or their combinations. For this reason the experimental procedure of [24] is extended here for the two methods tested therein, i.e. Passive Global Localisation-Fourier-Mellin Invariant matching with Centroids for translation (PGL-FMIC) and- Point-to-Line ICP (PGL-PLICP), and then for ALS, Monte Carlo Localisation (MCL) [9], and General Monte Carlo Localisation

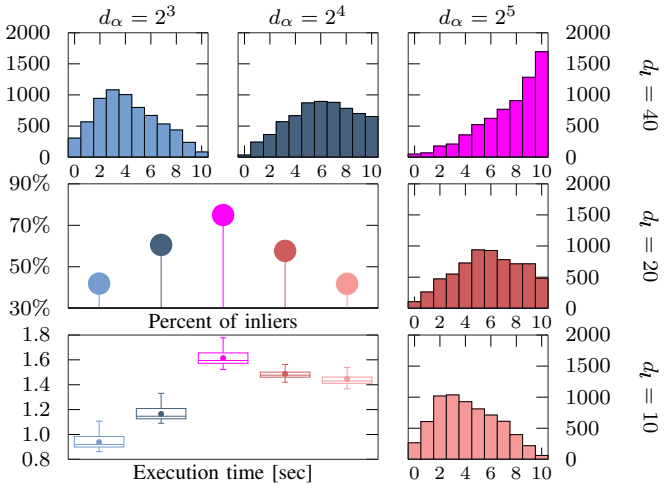

Fig. 5: Top row and right column: histograms of the number of times out of $N_E = 6669$ attempts of CBGL at global localisation when exactly $n \in [0, k] = [0, 10]$ pose estimates belonging to set $\mathcal{H}_1$ exhibited pose errors lower than $\delta_0 = 0.5$ (m^2 + rad^2)$^{1/2}$ (e.g. for $(\delta_l, \delta_\theta) = (0.3, 0.4)$[m, rad]). For densities $(d_l, d_\alpha) = (40, 2^5)$ (top right) this number is strictly increasing with n, which would be an aspect to be retrospectively expected of a global localisation method. Denser configurations adhere to the same pattern as well but would require more execution resources. Middle block: the resulting percent proportion of pose estimates whose pose error is lower than δ_0 for varying pose densities. Bottom block: the resulting distribution of corresponding execution times

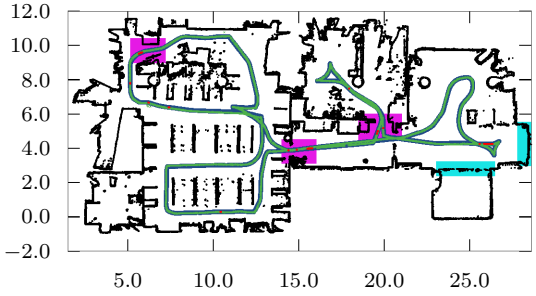

Fig. 6: The map of the real environment CSAL (black), the trajectory of the sensor (blue), and CBGL's estimated positions of the sensor (green). A total of $N_E = 6669$ pose estimations take place, 99.1% of which result in position errors lower than $\delta_l = 0.5$ m. See fig. 7 for the respective percentages of position and orientation errors for varying admissibility thresholds δ_l and δ_θ. Estimation is performed for each sensor pose independently of previous estimates and measurements. Sensor poses for which CBGL's output exhibits position error larger than $\delta_l = 0.5$ m are marked with red, sources of great range noise with cyan, and regions around doors with purple

(GMCL) [12]. All methods are tested against the two most challenging environments, i.e. WAREHOUSE (fig. 1; top) and WILLOWGARAGE, whose maps are Occupancy Grids, and in which a panoramic range sensor is placed at 16 different poses, respectively $\boldsymbol{p}_*^W$ and $\boldsymbol{p}_*^G$, for $N = 100$ independent attempts at global localisation per pose. The tests are conducted with the use of a sensor whose number of rays $N_s = 360$, maximum range $r_{\max} = 10.0$ m, and noise $\sim \mathcal{N}(0, 0.05^2)$ [m,m^2]. CBGL's internal `sm2` method is chosen to be `x1` [26] due to the periodicity of the range signal, `x1`'s robust pose errors compared to `sm2` state of the art methods, and its greater ability in matching

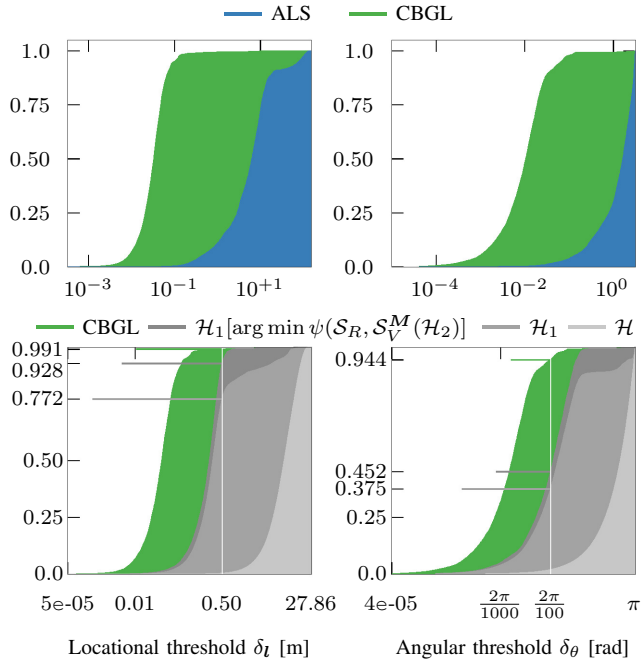

Fig. 7: Proportions of pose estimates whose position and orientation error is lower than corresponding thresholds δ_l and δ_θ. Top: ALS vs CBGL. Bottom: CBGL and internal pose estimate sets. Approximately 77% of all bottom-k pose estimates—the contents of $\mathcal{H}_1$ sets—exhibit position errors lower than $\delta_l = 0.5$ m for $k = 10$, and so do over 99% of CBGL's output pose estimates. The improvement in position and orientation induced by scan–to–map-scan matching is captured by the difference between the output (i.e. $\mathcal{H}_2[\arg\min\psi(\mathcal{S}_R,\mathcal{S}_V^M(\mathcal{H}_2)])$ and $\mathcal{H}_1[\arg\min\psi(\mathcal{S}_R,\mathcal{S}_V^M(\mathcal{H}_2)]$

scans captured from greater initial displacements than ICP alternatives [26]. The latter translates to the need for smaller initial hypotheses sets: for each environment the locational density is set to 3e+04 divided by the free space area of each environment. For Monte Carlo approaches MCL and GMCL the number of initial hypotheses is also set to 3e+04.

The maximum range of the sensor is such that the geometry of environment WAREHOUSE causes disorderly and extended lack of sampling of the sensor's surrounding environment, which limits available information and may therefore produce spurious measurements and increase ambiguities between the ranks of candidate estimates. In WILLOWGARAGE, on the other hand, almost all sensor placements result in complete sampling of its surroundings, but the sensor is purposefully posed in such conditions as to challenge the localisation methods' ability to perform fine distinctions between similar surroundings. Figure 8a depicts the overall distribution of position errors, orientation errors, and execution times per tested environment and algorithm, and fig. 8b depicts the percentage of outputs whose position error is lower than $\delta_l = 0.5$ m per tested pose. Although the cardinality of set $\mathcal{H}$ is equal in both environments, CBGL's execution times are uneven due to `x1`'s increased execution time when dealing with scans with missing range information. ALS is more robust against missing information than against repeated surroundings, while PGL-PLICP exhibits the inverse tendency. Despite the aforementioned sources of

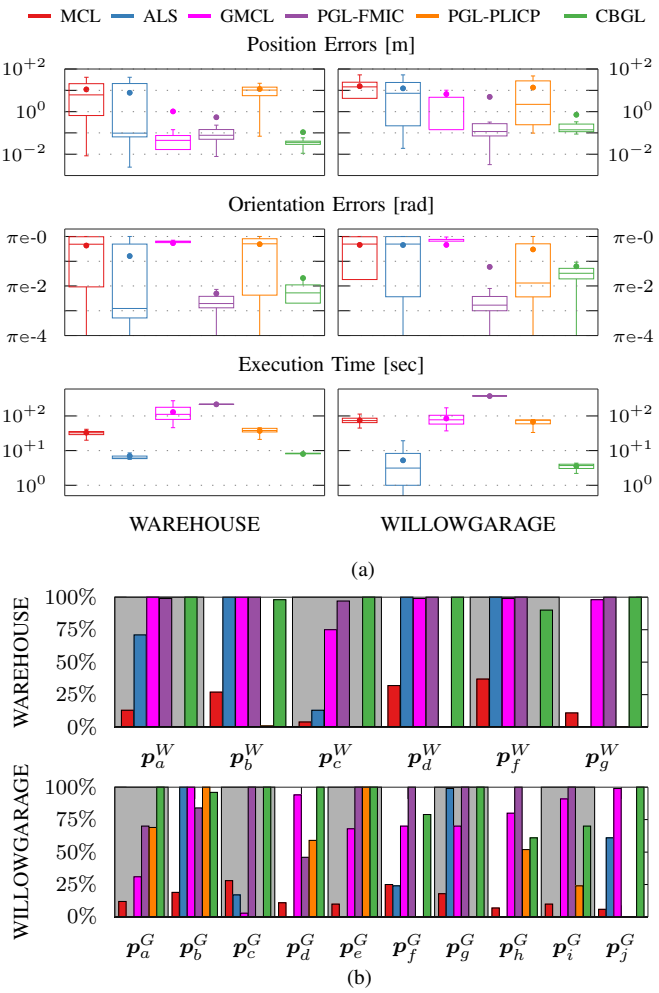

Fig. 8: (a) Distribution of position errors, orientation errors, and execution times per tested environment and algorithm in seconds for $N_s = 360$ rays. CBGL's execution time is at least eighteen times lower than other Monte Carlo approaches in WILLOWGARAGE and four times lower in WAREHOUSE. (b) Percent proportions of pose outputs whose position error is lower than $\delta_l = 0.5$ m per tested environment, pose, and method. Overall CBGL (green) features the highest number of inlier poses.

uncertainty, CBGL manages to exhibit the overall highest number of poses whose position error is below $\delta_l = 0.5$ m and the lowest mean position error.

C. Simulations against environmental and `sm2` algorithmic disparity

The third type of test aims to inquire how the performance of CBGL scales with respect to increasing environment area (and therefore increased number of hypotheses), environment diversity, sensor angular range, and choice of overlying `sm2` method. CBGL is tested once in each of $N_E = 45402$ environments, generated via the evaluation procedure of [26], which utilises five established and publicly available benchmark datasets provided courtesy of the Department of Computer Science, University of Freiburg [32]. The environments' sizes vary according to fig. 10 (left). Each coordinate of the Point Cloud map of each environment is corrupted by noise $\sim \mathcal{N}(0, 0.05^2)$ [m, m^2]. The angular range of the

range sensor varies according to the overlying `sm2` method used: for the Normal Distribution Transform (NDT) [33], Fast Generalised ICP (FastGICP) [34], and Fast Voxelised Generalised ICP (FastVGICP) [35]: $\lambda = 3\pi/2$ rad; for `x1`: $\lambda = 2\pi$ rad. Measurement noise is $\sim \mathcal{N}(0, 0.03^2)$ [m,m^2].

Figure 9 illustrates that, with the exception of NDT, all versions of CBGL exhibit mean positional errors less than 1.0 m; CBGL's combination with `x1` exhibits a mean error of approximately 0.5 m. The evidence illustrate that CBGL is robust to sensor angular range, as the distributions of errors between bottom-k ($\mathcal{H}_1$) sets are virtually indistinguishable for $k = 10$. Figure 10 (left) shows the execution time of CBGL combined with FastVGICP or `x1` as a function of environmental area, and the timing breakdown of the combination of CBGL with `x1` (right) with respect to (a) CBGL's total time minus `sm2` time and (b) computing map-scans, as proportions of total execution time.

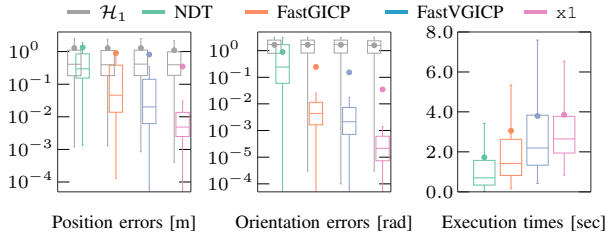

Fig. 9: Distributions of positional and orientational errors and of execution time of CBGL for varying choices of scan–to–map-scan matching methods. The errors of CBGL's internal $\mathcal{H}_1$ set are virtually unaffected by the decrease in angular range λ ($\lambda_{\text{NDT}} = \lambda_{\text{FastGICP}} = \lambda_{\text{FastVGICP}} = 3\pi/2 \neq \lambda_{\text{x1}} = 2\pi$)

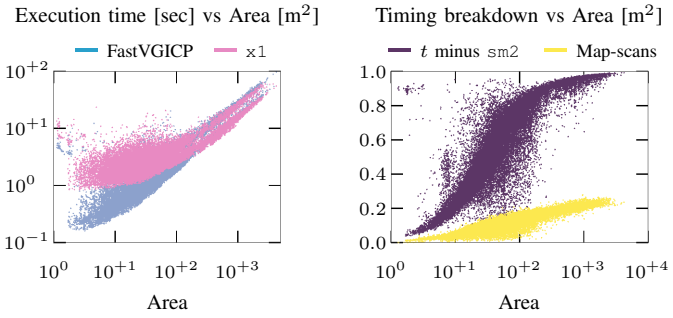

Fig. 10: Left: CBGL's execution time with respect to environment area for two choices of overlying scan–to–map-scan matching methods. In rough terms μ_t^{CBGL} [sec] $\simeq 10^{-2} \cdot area$ [m^2] for areas larger than 200 m^2. Right: Proportion of CBGL $\circ$ `x1`'s total execution time spent on (a) all operations up to and except for matching, and (b) computing map-scans, with respect to area

VI. Characterisation & Limitations

Range scans with panoramic angular ranges induce fewer pose ambiguities in pose estimate rankings than those with non-panoramic field of view. In the latter case this means that, given the evidence of subsections V-A and V-C where $\lambda = 3\pi/2$ rad, the choice of $k = 10$ largely inhibits the propagation of ambiguities to the output (fig. 9). However, non-panoramic sensors coupled with repeated environment structures may give rise to the conditions of fig. 11 (bottom). Other sources of potential, large pose errors for CBGL are portrayed in fig. 6: Regions coloured with cyan indicate closed glass doors, wherein high range errors in $\mathcal{S}_R$—which result from premature and arbitrary reflections of the LIDAR's light on glass—form major discrepancies with regard to map-derived virtual ranges. These subsequently propagate to CAER values and ultimately corrupt the estimation of pose error hierarchies derived from these values. Discrepancies of this kind arise also in regions coloured purple, which indicate vicinities around doors. In these areas the non-linear contour of ranges, combined with the fact that from different position estimates different parts of the environment become visible (and therefore small changes in position may result in large discrepancies between real and virtual scans)—these factors may cause to require higher values of locational density d_l or values of k in order to suppress highly erroneous pose estimates propagated to the output.

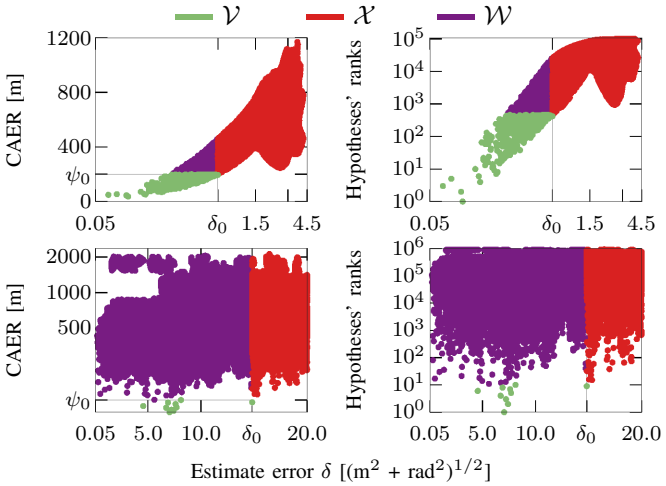

Fig. 11: Top figures: In principle there is at least one admissible pose estimate ($\delta \leq \delta_0$) in $\mathcal{H}$ for any choice of k. However, k sets threshold ψ_0 on the CAER of estimates in $\mathcal{H}$, and therefore it sets $\mathcal{V}$, where $\mathcal{V} \cup \mathcal{X} \cup \mathcal{W} = \mathcal{H}$. In case of (a) repetitions of the immediate environment of the sensor in a given map, and (b) non-panoramic angular range of a sensor, it is possible that a choice of k may starve $\mathcal{V}$ of admissible poses, as witnessed in the two figures at the bottom (this is also true for pose $\boldsymbol{p}_i^G$ in environment WILLOWGARAGE of subsection V-B). See [27] for details on $\mathcal{V}$, $\mathcal{X}$, and $\mathcal{W}$

VII. Conclusions and Future Steps

This article has presented a single-shot Monte Carlo approach to the solution of the passive version of the global localisation problem with the use of a 2D LIDAR sensor, titled CBGL. CBGL allows for the fast estimation of the sensor's pose within a metric map by first dispersing hypotheses in it and then leveraging (a) the proportionality of values of the Cumulative Absolute Error per Ray (CAER) metric to the position and orientation errors of the hypotheses, for estimates in a neighbourhood of the sensor's pose, and (b) the lack of disproportionality outside of it. CBGL was evaluated in various real and simulated conditions and environments; it was found to be superior to Monte Carlo and feature-based approaches in terms of number of inlier pose estimates and

execution time. Future steps will aim at (a) the extension of the CAER metric for the use with 3D LIDAR sensors in service to a solution of the problem of their localisation in 6DoF, and (b) making CBGL more robust by considering the statistics of clusters in case of disparate estimates in its $\mathcal{H}_2$ set. The C++ ROS code of the proposed method is available at `https://github.com/li9i/cbgl`.

REFERENCES

[1] A. Millane, H. Oleynikova, J. Nieto, R. Siegwart, and C. Cadena, "Free-Space Features: Global Localization in 2D Laser SLAM Using Distance Function Maps," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, nov 2019, pp. 1271–1277.

[2] K. P. Cop, P. V. K. Borges, and R. Dube, "Delight: An Efficient Descriptor for Global Localisation Using LiDAR Intensities," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2018, pp. 3653–3660.

[3] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, J. Wan, and N. M. Kwok, "A Comprehensive Performance Evaluation of 3D Local Feature Descriptors," *International Journal of Computer Vision*, vol. 116, no. 1, pp. 66–89, jan 2016.

[4] A. Tzitzis, A. Filotheou, A. R. Chatzistefanou, T. Yioultsis, and A. G. Dimitriou, "Real-Time Global Localization of a Mobile Robot by Exploiting RFID Technology," *IEEE Journal of Radio Frequency Identification*, pp. 1–1, 2023.

[5] S. Lowry, N. Sünderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford, "Visual Place Recognition: A Survey," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 1–19, feb 2016.

[6] D. M. Rosen, K. J. Doherty, A. Terán Espinoza, and J. J. Leonard, "Advances in Inference and Representation for Simultaneous Localization and Mapping," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, no. 1, pp. 215–242, may 2021.

[7] M. Cooper, J. Raquet, and R. Patton, "Range Information Characterization of the Hokuyo UST-20LX LIDAR Sensor," *Photonics*, 2018.

[8] H. Yin, X. Xu, S. Lu, X. Chen, R. Xiong, S. Shen, C. Stachniss, and Y. Wang, "A survey on global lidar localization: Challenges, advances and open problems," *arXiv:2302.07433*, 2023.

[9] N. Akai, L. Y. Morales, and H. Murase, "Mobile Robot Localization Considering Class of Sensor Observations," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, oct 2018, pp. 3159–3166.

[10] J. Wang, P. Wang, and Z. Chen, "A novel qualitative motion model based probabilistic indoor global localization method," *Information Sciences*, vol. 429, pp. 284–295, mar 2018.

[11] A. Yilmaz and H. Temeltas, "Self-adaptive Monte Carlo method for indoor localization of smart AGVs using LIDAR data," *Robotics and Autonomous Systems*, vol. 122, p. 103285, dec 2019.

[12] M. A. Alshikh Khalil and I. Hatem, "Development of a new technique in ros for mobile robots localization in known-based 2d environments," *Tishreen University Journal of Research and Scientific Studies - Engineering Sciences Series*, pp. 119–137, 2021.

[13] R. Chen, H. Yin, Y. Jiao, G. Dissanayake, Y. Wang, and R. Xiong, "Deep Samplable Observation Model for Global Localization and Kidnapping," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2296–2303, apr 2021.

[14] F. Kallasi and D. L. Rizzini, "Efficient loop closure based on FALKO lidar features for online robot localization and mapping," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, oct 2016, pp. 1206–1213.

[15] M. Usman, A. M. Khan, A. Ali, S. Yaqub, K. M. Zuhaib, J. Y. Lee, and C.-S. Han, "An Extensive Approach to Features Detection and Description for 2-D Range Data Using Active B-splines," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2934–2941, jul 2019.

[16] Z. Wang, L. Zhang, S. Zhao, and S. Zhang, "Global Localization With a Single-Line LiDAR by Dense 2D Signature and 1D Registration," *IEEE Sensors Journal*, vol. 21, no. 10, pp. 11 497–11 506, may 2021.

[17] J. Meng, L. Wan, S. Wang, L. Jiang, G. Li, L. Wu, and Y. Xie, "Efficient and Reliable LiDAR-Based Global Localization of Mobile Robots Using Multiscale/Resolution Maps," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–15, 2021.

[18] R. W. M. Hendrikx, P. Pauwels, E. Torta, H. J. Bruyninckx, and M. J. G. van de Molengraft, "Connecting Semantic Building Information Models and Robotics: An application to 2D LiDAR-based localization," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2021, pp. 11 654–11 660.

[19] S.-Y. An and J. Kim, "Extracting Statistical Signatures of Geometry and Structure in 2D Occupancy Grid Maps for Global Localization," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4291–4298, apr 2022.

[20] K. Nielsen and G. Hendeby, "Survey on 2D Lidar Feature Extraction for Underground Mine Usage," *IEEE Transactions on Automation Science and Engineering*, vol. 20, no. 2, pp. 981–994, apr 2023.

[21] T.-X. Xu, Y.-C. Guo, Z. Li, G. Yu, Y.-K. Lai, and S.-H. Zhang, "TransLoc3D : Point Cloud based Large-scale Place Recognition using Adaptive Receptive Fields," may 2021. [Online]. Available: http://arxiv.org/abs/2105.11605

[22] P. Yin, F. Wang, A. Egorov, J. Hou, Z. Jia, and J. Han, "Fast Sequence-Matching Enhanced Viewpoint-Invariant 3-D Place Recognition," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 2, pp. 2127–2135, feb 2022.

[23] J. Komorowski, "Improving Point Cloud Based Place Recognition with Ranking-based Loss and Large Batch Training," in *2022 26th International Conference on Pattern Recognition (ICPR)*. IEEE, aug 2022, pp. 3699–3705.

[24] A. Filotheou, A. Tzitzis, E. Tsardoulias, A. Dimitriou, A. Symeonidis, G. Sergiadis, and L. Petrou, "Passive Global Localisation of Mobile Robot via 2D Fourier-Mellin Invariant Matching," *Journal of Intelligent & Robotic Systems*, p. 26, 2022.

[25] A. Filotheou, G. D. Sergiadis, and A. G. Dimitriou, "FSM: Correspondenceless scan-matching of panoramic 2D range scans," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022.

[26] A. Filotheou, A. L. Symeonidis, G. D. Sergiadis, and A. G. Dimitriou, "Correspondenceless scan-to-map-scan matching of 2D panoramic range scans," *Array*, 2023.

[27] A. Filotheou, "Cbgl: Fast monte carlo passive global localisation of 2d lidar sensor—mathematically-rigorous version," *arXiv:2307.14247v3*, 2024. [Online]. Available: https://arxiv.org/pdf/2307.14247v3.pdf

[28] G. Vasiljević, D. Miklić, I. Draganjac, Z. Kovačić, and P. Lista, "High-accuracy vehicle localization for autonomous warehousing," *Robotics and Computer-Integrated Manufacturing*, 2016.

[29] I. Vizzo, T. Guadagnino, B. Mersch, L. Wiesmann, J. Behley, and C. Stachniss, "KISS-ICP: In Defense of Point-to-Point ICP – Simple, Accurate, and Robust Registration If Done the Right Way," *IEEE Robotics and Automation Letters*, 2023.

[30] N. Akai, "Reliable monte carlo localization for mobile robots," *arXiv:2205.04769*, 2022.

[31] A. Censi, "An ICP variant using a point-to-line metric," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, may 2008, pp. 19–25.

[32] "Datasets used in the experimental procedure." [Online]. Available: http://ais.informatik.uni-freiburg.de/slamevaluation/datasets.php

[33] P. Biber and W. Strasser, "The normal distributions transform: a new approach to laser scan matching," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, vol. 3. IEEE, 2003, pp. 2743–2748.

[34] A. Segal, D. Haehnel, and S. Thrun, "Generalized-ICP," in *Robotics: Science and Systems V*. Robotics: Science and Systems Foundation, jun 2009.

[35] K. Koide, M. Yokozuka, S. Oishi, and A. Banno, "Voxelized GICP for Fast and Accurate 3D Point Cloud Registration," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2021, pp. 11 054–11 059.