Geometry matters: Place Recognition in 2D Range Scans using Geometrical Surface Relations

Marian Himstedt, Erik Maehle

***Abstract*— Place recognition is a fundamental requirement for mobile robots. It is particularly needed for detecting loop closures in SLAM and to enable self-localization for mobile robots given a prior map. The multitude of existing approaches rely on appearance based methods, e.g. the extraction of interest points in terms of local extrema. It can be observed that the availability of these features is highly environment specific and the limited descriptiveness causes a large number of false-positive matches. This paper utilizes a generic environment description based on normal surface primitives. The association of different places is done using Geometrical Surface Relations (GSR) of co-occurring primitives. Experimental results obtained from publicly available datasets demonstrate that GSR outperforms state-of-the-art approaches in place recognition for large scale outdoor as well as indoor environments.**

I. Introduction

Place recognition is an essential capability of a mobile robot in order to fulfill high-level tasks. In particular, it is necessary for detecting loop closures in SLAM as well as the global localization with respect to a prior map. In addition to these common applications, place recognition contributes to a number of further tasks, as for instance the merging of datasets collected during several surveys. Places are associated either in a pure topological or metric manner. The latter estimates a relative transformation of the query scan with respect to the matching one, whereas appearance based approaches solely match nodes of topological graphs **[1].** While vision-based solutions to place recognition have been quite exhaustively investigated, only little attention has been paid to the capabilities of 2D range data. This is obviously due to the fact that, in contrast to images, 2D range scans only describe a relatively small subspace of the surrounding environment and hence provide less information to distinguish different places. Exhaustive matching of invidual laser scans is impractical in terms of time complexity when looking at large scale environments making descriptive place signatures and efficient retrieval algorithms indispensable. Similarly to vision-based methods, the extraction of local interest points serving as landmarks is beneficial. In existing approaches these landmarks are assigned descriptors capturing their local surroundings. As usual in bag-of-words (BOW) techniques these descriptors can be quantized to words enabling laser scans to be represented by histograms of word occurrences. While achieving promising results in visual applications [1], BOW performs rather poor in conjunction with 2D range data [2], particulary when looking at outdoor environments. Inspired by the vision community [3], Tipaldi et al. introduced Geometrical FLIRT phrases (GFP) allowing approximate geometric verification by preserving the clock-wise order of observed words [2]. Although GFP reduces the number of potential place correspondences, the number of required full geometric verifications on outdoor sequences increases with the size of the environment. This paper introduces Geometrical Surface Relations (GSR) explicitly taking into account relative metric displacements of landmarks observed in individual range scans.

Marian Himstedt and Erik Maehle are with the Institute of Computer Engineering, University of Lübeck, Germany.
E-mail: himstedt@iti.uni-luebeck.de

II. Related work

This section presents an overview of existing work dealing with place recognition. First, it is shown how GSR is related to vision-based approaches utilizing geometric constraints. Subsequently our approach is justified against prior work in range based place recognition. Two approaches being mostly related to our work are outlined in detail.

A. Spatial relations in vision-based place recognition

There exists a large amount of work dedicated to place recognition in the computer vision literature. A survey of place recognition in visual SLAM, i.e. loop closure detection, is given by [4]. Matching of geometric constraints within visual recognition tasks has recently drawn potential interest again. In [5] Johns and Yang propose to quantize the relative orientations and distances of co-occurring features in the image coordinate frame. Also Cho et al. incorporate spatial relationships between keypoints in a graph-based object recognition system. Paul and Newman extended FABMAP, an apperance based place recognition system, by modelling spatial relations of observed keypoints in random graph models [6]. In contrast to [5], FABMAP-3D does not work in the image coordinate frame.

Clemente et al. propose to use relative distances of observed features as an evaluation step when matching submaps [7]. However, the complexity of the matching procedure limits its use for rather small numbers of submaps. Finman et al. proposed physical words to encode spatial relations of objects segmented from RGB-D images, captured solely in indoor environments [8].

B. Place recognition based on range scans

The related work on place recognition utilizing 2D lidar scans is rather small compared to those using camera images.

Bosse and Zlot analyse several detectors and descriptors for local keypoints aggregated from sequential 2D range scans into submaps. The recognition of the submaps is carried out by associating stored and observed feature descriptors using approximate nearest neighbour (ANN) search with each correspondence voting for the associated submap [9]. The submaps having the highest numbers of votes are selected and checked for geometric consistency based on projection histograms and ICP.

Granström et al. concatenate a set of various features such as curvature and average range serving as input for an Adaboost based classifier that detects corresponding scans [10]. As shown in [2] this approach is computationally very expensive compared to the state-of-the-art.

All of the mentioned approaches [9], [2], [10] extract features from raw data serving as keypoints. Different from that, authors utilized occupancy grid maps allowing to incorporate submaps built incrementally and potentially using multiple sensors. In our recent work we investigated non-negative matrix factorization (NMF) to automatically determine relevant feature types of a given dataset in terms of basis primitives constructed from adjacent grid cells [11]. By representing local grid maps as a distribution of basis primitives, the presented method also allows to reduce the storage sizes of the prior occupancy grid maps which scale quadratically with increasing environment sizes.

In addition to presented work there is also a number of approaches to place recognition working on 3D range data. Steder et al. introduced a recognition system which extracts points of local extrema and uses small patches of the range image around the latter [12]. Similarly to [2] the features are quantized and matched using a BOW model. Magnusson uses normal surface primitives whose appearances (e.g. shape) are classified. Feature vectors encoding the frequencies of present surface primitives are used to identify places [13]. Our approach shares the idea of using surface primitives, but passes on modelling the appearances.

C. Geometrical FLIRT Phrases

The most related work to our approach is given by [2], [14] and hence specifically outlined here. Tipaldi and Arras presented FLIRT, a multi-scale feature detector for 2D range scans [14]. This work was accompanied by a shape context and a beta grid desciptor. The latter comprises raw distance measurements around local interest points into polar tesselations which were shown to be superior to the well-known shape context descriptors. Tipaldi et al. extended their approach by encoding the local descriptors as words which are assigned according to an a-priori learnt vocabulary [2]. While pure appearance based place signatures based on bag-of-words (BOW) achieve rather poor results, particularly for outdoor datasets, the authors introduced Geometrical FLIRT phrases (GFP). Inspired by the visual recognition system of Zhang et al. [3], GFPs enable to match sets of features in the cyclic order of observation ensuring the preservation of geometric constraints. GFP performs better than BOW but still requires a large number of full geometric consistency checks for potential matches in order to achieve competitive recognition results.

D. Geometrical Landmark Relations

We recently presented Geometrical Landmark Relations (GLARE) [15] which similarly to FLIRT extracts local features of high curvature from range scans. However, it does not generate local descriptors to identify the features but instead makes use of geometrical relations of co-occurring features in a range scan. We demonstrated that aggregating relative distances and orientations of co-occurring landmarks to scan signatures enables us to match place correspondences among very large datasets with significantly less geometrical consistency checks. GLARE was shown to outperform GFP in run time and recognition performance for outdoor environments. GFP however achieved slightly better recognition rates in indoor environments.

E. Summary

The majority of the presented work uses local interest points with quantized descriptors and requires vocabularies or prior training stages. Except for GLARE [15], none of the presented approaches working on range scans makes explicit use of relative orientations and distances of co-occurring landmarks which can substantially contribute towards distinguishing places.

III. Geometrical Surface Relations

This section introduces Geometrical Surface Relations (GSR). At first it is explained how GSRs are generated from the input 2D range scans. Subsequently it is shown how these signatures can be integrated into efficient retrieval algorithms.

A. Extraction of surface primitives

At first, all measurements of a 2D range scan are projected onto a regular grid. Similar to Magnusson et al. [13] we estimate a normal distribution with mean μ_i and a covariance matrix Σ_i for the measurements of each non-empty cell i. The surface orientations θ_i of each primitive is required in order to model the spatial relations. For this purpose we utilize an eigenvalue decomposition of the surface primitive's covariance matrix. The eigenvector e_{min} with the smallest eigenvalue is selected for estimating the orientation $\hat{\theta}_i$:

$$\hat{\theta}_i = atan2(e_{min}^{\{y\}}, e_{min}^{\{x\}}) \tag{1}$$

Since the eigenvector e_{min} is not necessarily pointing towards the sensor's origin, we explicitely account for this by estimating the primitives' orientation ψ_{orig} towards the sensor's origin:

$$\psi_{orig} = atan2(\mu_i^{\{y\}}, \mu_i^{\{x\}}) + \pi \tag{2}$$

If the displacement of ψ_{orig} and $\hat{\theta}_i$ exceeds a threshold τ_{max}, we map the surface primitive's orientation as follows:

$$\theta_i = \begin{cases} \hat{\theta}_i & if\,(\psi_{orig} - \hat{\theta}_i) < \tau_{max} \\ \hat{\theta}_i + \pi & otherwise \end{cases} \tag{3}$$

In this way it is ensured that θ_i is assigned the expected direction. We empirically set $\tau_{max} = \pi/3$, which however is not too crucial. The size and resolution of the grid have to be justified according to the type of environment and the sensor used. As a result of this step we obtain a set of surface primitives $l_i = \{\mu, \Sigma, \theta\}_i$ for each range scan.

B. Encoding spatial relations

For a set of N surface primitives $l_1, \ldots, l_N$ detected in the k-th range scan we estimate the Euclidean distances $\rho_{i,j}$ of each primitive $l_i = \{\mu, \Sigma, \theta\}_i$ to all others $l_j = \{\mu, \Sigma, \theta\}_j$ of the set with $i \neq j$ within the local coordinate frame of the range scan. In addition to that the bearings $\Delta\theta_{i,j}$ and $\Delta\theta_{j,i}$ of co-occurring primitives are estimated:

$$\Delta\theta_{i,j} = \theta_i - \theta_j \tag{4}$$

Having estimated $\rho_{i,j}$ and $\Delta\theta_{i,j}$, the distribution of spatial relations is captured. For this purpose the bearings $\Delta\theta_{i,j}$ and distances $\rho_{i,j}$ are quantized and assigned to uniformly sized bins:

$$(\Delta\theta_{i,j}, \rho_{i,j}) \in bin\,(n_\theta, n_\rho) \tag{5}$$

The 2D histogram $H_{i,j}$ is constructed by a discrete multivariate Gaussian centred on the bin $n = (n_\rho, n_\theta)$ given the covariance matrix Σ_H. Centering the Gaussian to the discretized bin n instead of the exact position $(\Delta\theta_{i,j}, \rho_{i,j})$ allows a precomputation of the contributions to the histogram bins covered by the covariance Σ_H. The histogram position $m = (m_\rho, m_\theta)$ is estimated according to:

$$H_{i,j}\,(m) = \mathcal{N}\,(m - n,\, \Sigma_H) \tag{6}$$

The signature $S^{(k_i)}$ for the primitive l_i is estimated as follows:

$$S^{(k_i)} = \sum_j H_{i,j} \tag{7}$$

Finally, we obtain the scan signature $S^{(k)}$ for the range scan k by estimating a normalized sum over all signatures:

$$S^{(k)} = \eta \sum_i S^{(k_i)} \tag{8}$$

The factor η ensures the normalization of the signature. Figure 1 exemplarily illustrates a GSR signature.

Our technique differs from a multitude of other place recognition approaches since the latter typically quantize the local neighbourhood of features to descriptors which are quantized to words of a prior vocabulary [14], [5]. Each descriptor hence votes for exactly one word which is also referred to as hard-voting. In contrast, we use a soft-voting scheme enabling to evaluate adjacent cells depending on Σ_H. The covariance matrix Σ_H allows to incorporate uncertainties in relative depth and orientation estimates. These are mainly driven by the sensor accuracy but can also be biased by the localization quality of the extracted features. We observed that this scheme results in smoother results which is particularly well-suited for noisy range data.

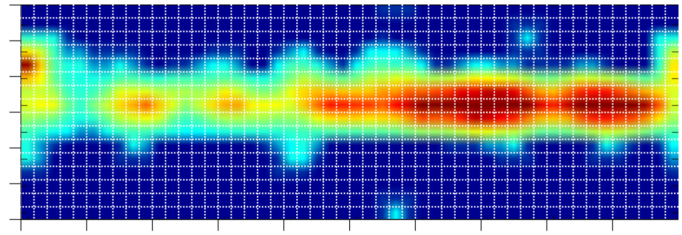

Figure 1. Example of a scan signature generated using Geometrical Surface Relations.

C. Efficient scan retrieval

Having generated the scan signatures, the goal is to collect those in a global repository S while simultaneously managing an index structure allowing fast scan retrieval. Since S consists of high-dimensional vectors, Approximate Nearest Neigbor Search (ANN) becomes indispensable in order to enable efficient retrieval [16]. Conventional kd-trees always pick that dimension of the input data having the largest variance to bisect the data at each level of the tree. This structure has shown to be less suitable in the presence of high-dimensional data [17]. Thus it is recommended to make use of multiple randomized kd-trees whose splitting bounds are randomly drawn from the top variant dimensions resulting in better representations for higher dimensions. The key difference of ANN compared to the exact nearest neighbour search is that the number of points to evaluate is bound to τ. This method achieves a sufficient approximation providing the bound is set appropriately. Setting this bound is critical as it poses a trade-off of accuracy and time complexity. In addition to that, it is affected by the balance quality of the kd-tree. When using GSR for recognizing places with respect to a prior map (global localization) we can assume to have a well-balanced tree. For loop closure detection in SLAM, the global repository S is constantly updated while traversing the environment which requires the underlying kd-trees to be rebalanced after a certain amount of change. We found that incorporating the ratio of the kd-tree leafs and the number of recently appended elements (before rebalancing) indicate a promising choice of τ.

Given the scan signature $S^{(k)}$, the global repository S is searched for potential candidate matches using a cosine distance function. For performance reasons we limit the search to the K approximate nearest neighbours. Similiarly to τ, K needs to be set appropriately in order to obtain optimal results in terms of precision and recall. Setting K to a lower value might result in missing place correspondences, whereas large values of K require a potentially large number of geometric consistency checks which are computationally expensive due to the trigonometric operations involved.

We observed that the recognition of GSR based signatures using ANN achieves optimal results in terms of time complexity and accuracy. The inverted file techniques that are commonly used for vision-based place recognition [5], [1] and also in GFP [2] are not suited for GSR due to their dense histogram representations. The distributions of visual words for a single location are typically rather sparse representing only a small subset of the prior vocabulary. Combined with a hard-voting scheme this allows to invert the recognition by having visual words pointing towards their corresponding images or scans. Only those places being drawn are considered for further matching. Applied to GSR, this procedure does not contribute to a faster retrieval since on the one hand the number of votes in the histogram is significantly larger (N vs. $N^2/2$). Second, our soft-voting mechanism does not solely consider the closest bin but also adjacent ones.

D. Geometric verification

Due to *perceptual aliasing* it is inevitable that two different places might be close in the signature space. This particularly occurs in environments with a higher number of repetitive structures such as long corridors in indoor environments. That is why the set of K putative matches for the place $S^{(k)}$, provided by the retrieval system, is checked for geometrical consistency. For this purpose, we again make use of the geometric relations of co-occurring landmarks. For each surface primitive l_i in the k-th scan, the landmark signature $S^{(k_i)}$ is estimated according to Eq. 7. The landmark correspondences of place $S^{(k)}$ to the K putative matches are obtained by estimating the pair-wise distances of the landmark signatures $S^{(k_i)}$. Those being above a fixed threshold are considered for a RANSAC based scan-matching which estimates a rigid transformation of the two scans. The residual error as well as the number of inlier correspondences indicate whether the two given scans satisfy our geometric consistency check and hence are considered as matching places.

IV. EXPERIMENTS

In order to evaluate the presented approach, a number of experiments were carried out. The first experiment shows the recognition performance on four publicly available datasets. Here the goal is to compare the results of Geometrical Surface Relations (GSR) to the state-of-the-art approaches GLARE [15] and GFP [2]. The second experiment analyzes the recognition performance of GSR over longer periods of time. While the first experiment is rather focussing on place recognition for detecting loop closures in SLAM (whole trajectory considered), the second rather demonstrates the GSRs performance for global localization and relocalization respectively.

A. GSR vs. State-of-the-art

For this experiment we selected four different publicly available datasets of indoor and outdoor environments which were also used in [15] and [2] to evaluate GLARE and GFP respectively (see Table I).

Table I
DATASETS FOR EXPERIMENTS

Name	Type	# Scans	# Features	Path length [m]
Intel-lab	indoor	2672	39392	360.7
FR-Clinic	outdoor	6917	190760	1437.6
Victoria Park	outdoor	5751	81795	4206.14
Kenmore	outdoor	13063	499237	6588.34

GFP are generated using the optimal settings, as shown in [2]. For GLARE and GSR we use 8 angular bins, 100 linearly sized bins with a size of 0.5 m for the outdoor datasets and 50 bins with a size of 0.2 m for the indoor datasets (see [15]). GSR is initialized with grid cell sizes of 0.25 m for indoor and 1.0 m for outdoor datasets. We tested GSR, GLARE and GFP on all datasets with a number of $K = 50$ nearest neighbours. The geometric verification rejects putative matches of places by thresholding based on the residual error (linear: $0.5m$, angular: $0.2rad$). The thresholds for estimating the precision/recall curves are obtained by using a different numbers of inlier correspondences c (here: $a = 1$, 2, ..., 32). The datasets are matched one-by-one, but ignoring the trivial self-match. This procedure is similar to the procedures in [2] and [15]. The results are shown in Figure 2.

B. Long-term Recognition using GSR

This experiment evaluates the recognition of GSR over longer periods of time. For this purpose we make use of a subset of five datasets of the MIT Stata Center collection [18] covering a total period of time of more than 2 months (see Table II). The ground truth supplemented with these datasets provides an accuracy of about 2cm. The GSRs generated for the first dataset (2012-01-18) serve as prior map for the subsequent datasets. This demonstrates the recognition performance of GSR in terms of global localization. The results are shown in Figure 3. Only the results of the first dataset (2012-01-18) are obtained similar to the first experiment by matching the dataset to itself. The ground truth poses of datasets 2-5 are used to find overlapping areas of the environment traversed. Hence we exclude those scans of the datasets 2-5 that are mapped outside the area of dataset 1. The settings for GSR are similar to the ones for the indoor environments of experiment 1.

Table II
STATA CENTER DATASETS FOR LONG-TERM EXPERIMENT.

#	Date	# Scans	# Primitives	Path length [m]
1	2012-01-18	2562	237559	683.0
2	2012-01-25	1355	135246	348.0
3	2012-01-28	2279	212455	635.0
4	2012-02-02	2806	275128	1003.0
5	2012-04-02	1726	154336	606.0

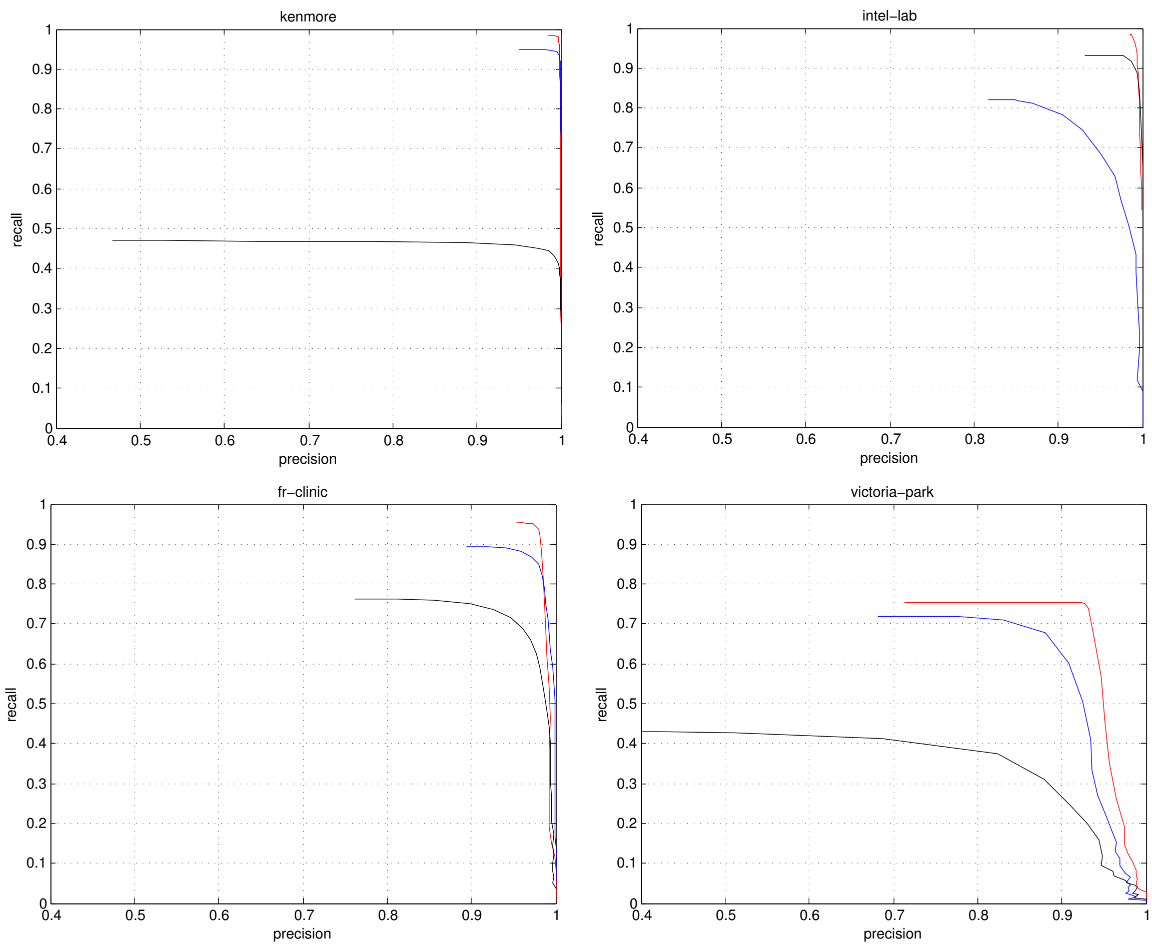

Figure 2. Experimental results obtained using GSR (red), GLARE (blue) and GFP (black) are shown for $K = 50$ nearest neighbours taken into consideration.

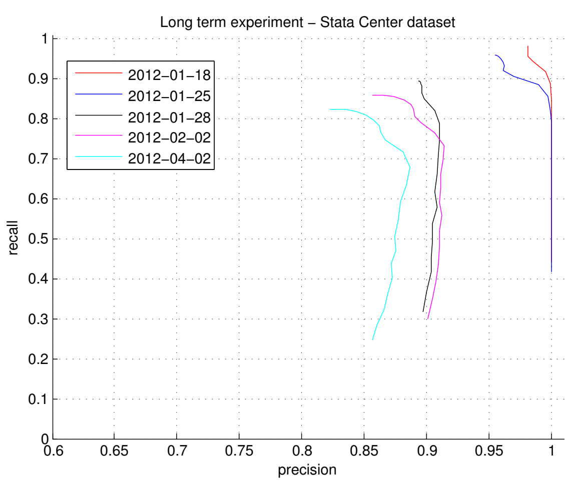

Figure 3. Place recognition results using GSR based on a subset of Stata Center dataset collection.

C. Run time

The run times for GSR, GLARE and GFP are shown in Table III. GFP only gets close to GLARE's and GSR's recall rates for a large number of nearest neigbors which is why $K = 500$ is used for this experiment. Actually GFP requires even more putative neighbours to be considered, however, the run time for this becomes highly impracticable.

The run times for feature detection for GLARE/GFP are very similar since they share the same algorithms (see Section III-A). The feature detection and description phase of GSR takes slightly longer than GLARE/GFP since generating the local normal surface primitive map is more time consuming. It is apparent that the geometric verification for GSR/GLARE is significantly faster than GFP. This is due to the fact that the relative distances of landmarks allow to reject more false-positive feature correspondences than the appearance based descriptors of GFP. We observed that the majority of features on the Kenmore dataset refer to point-like features resulting in very similar GFP descriptors. Thus the number of putative correspondences passed to RANSAC is smaller for GSR/GLARE compared to GFP. Note that GSR already achieved very accurate results for $K = 50$ (see Figure 2), which would substantially reduce the run time needed for geometric verification.

Table III
RUN TIME ON KENMORE FOR $K = 500$

	GSR [ms]	GLARE [ms]	GFP [ms]
feature detection	11.121	7.298	7.276
feature description	6.322	1.291	1.797
scan retrieval	26.198	25.312	81.462
geometric verification	102.596	88.928	330.586
total	**146.237**	**122.829**	**421.121**

D. Discussion

It is obvious that GSR outperforms both GFP and GLARE on all datasets in the first experiment. We observed that GFP gets closer to GLARE and GSR with a high number of K nearest neighbours taken into consideration which again confirms the results we achieved in prior experiments in [15]. The differences in precision/recall for GSR are rather small for increasing number of K. The number K essentially determines the run time of the place recognition since the geometric verfications that need to be carried out for K scans are computationally highly expensive. Even though GLARE already achieves promising results, the surface primitives of GSR further help distinguishing similar places, particularly in structureless environments.

The second experiment quantitatively shows GSRs performance for long-term place recognition. The recognition is worse than for the first experiment which is mainly due to the fact that the robot is likely to move further from the reference paths given by the first dataset. Secondly, it is obvious that recognition performance drops with increasing time difference to the reference dataset which is probably due to structural changes in the environment. This problem has to be tackled differently by taking changes into consideration. This problem, however, is not in the scope of this paper as we are rather interested in demonstrating the repeatability of our presented approach to place recognition.

V. CONCLUSIONS

This paper introduced Geometrical Surface Relations (GSR) designed for place recognition in 2D range data. GSRs encode relative landmark relations captured from single range scans into scan signatures. The latter reduces the scan retrieval to a very efficient approximate nearest neighbour search avoiding to evaluate each single signature of the repository for finding correspondences. It was shown that GSR outperforms the state-of-the-art approaches GLARE and GFP in outdoor as well as indoor environments. In contrast to GFP, GSR does not require any prior training stage emphasising its strengths for the application in a-priori unknown environments. Our experiments further demonstrated that GSR is very well suited for the global localization of a mobile robot. Similarly to the state-of-the-art GSR not only allows to recognize places but also to estimate a relative transformation and alignment error of the underlying range scans. This enables precise global localization and a convenient integration into existing graph based SLAM frameworks. Thanks to the more generic descriptive power of surface primitives, GSR provides an outstanding recognition performance for a multitude of environment types ranging from suburbs, parks, hallways, offices which was demonstrated on the datasets kenmore, victoria park, stata center and intel respectively. Thanks to the elliptic shapes of the surface primitives, GSR is able to automatically adapt to dominant features in these environments, as for instance trees, walls or doors.

REFERENCES

[1] M. Cummins and P. Newman, "FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance," *The International Journal of Robotics Research*, vol. 27, no. 6, pp. 647–665, 2008.

[2] G. D. Tipaldi, L. Spinello, and W. Burgard, "Geometrical flirt phrases for large scale place recognition in 2d range data." in *IEEE International Conference on Robotics and Automation (ICRA)*, Karlsruhe, Germany, 2013.

[3] Y. Zhang, Z. Jia, and T. Chen, "Image retrieval with geometry-preserving visual phrases," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 809–816.

[4] B. Williams, M. Cummins, J. Neira, P. Newman, I. Reid, and J. Tardós, "A comparison of loop closing techniques in monocular SLAM," *Robotics and Autonomous Systems*, 2009.

[5] E. Johns and G.-Z. Yang, "Dynamic scene models for incremental, long-term, appearance-based localisation," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2013, pp. 2731–2736.

[6] R. Paul and P. Newman, "Fab-map 3d: Topological mapping with spatial and visual appearance," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2010.

[7] L. A. Clemente, A. J. Davison, I. D. Reid, J. Neira, and J. D. Tardós, "Mapping large loops with a single hand-held camera," in *Robotics: Science and Systems*, 2007.

[8] R. Finman, T. Whelan, L. Paull, and J. J. Leonard, "Physical words for place recognition in dense RGB-D maps," in *ICRA workshop on visual place recognition in changing environments*, Jun. 2014.

[9] M. Bosse and R. Zlot, "Keypoint design and evaluation for place recognition in 2d lidar maps," *Robotics and Autonomous Systems*, vol. 57, no. 12, pp. 1211–1224, 2009.

[10] K. Granström, T. B. Schön, J. I. Nieto, and F. T. Ramos, "Learning to close loops from range data." *International Journal of Robotic Research*, vol. 30, no. 14, pp. 1728–1754, 2011.

[11] S. Hellbach, M. Himstedt, and H.-J. Boehme, "What's around me: Towards non-negative matrix factorization based localization," in *Mobile Robots (ECMR), 2013 European Conference on*, Sept 2013, pp. 228–233.

[12] B. Steder, G. Grisetti, and W. Burgard, "Robust place recognition for 3d range data based on point features," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2010.

[13] M. Magnusson, H. Andreasson, A. Nuchter, and A. Lilienthal, "Appearance-based loop detection from 3d laser data using the normal distributions transform," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2009.

[14] G. Tipaldi and K. Arras, "Flirt - interest regions for 2d range data," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 3616 –3622.

[15] M. Himstedt, J. Frost, S. Hellbach, H.-J. Boehme, and E. Maehle, "Large scale place recognition in 2d lidar scans using geometrical landmark relations," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept 2014.

[16] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu, "An optimal algorithm for approximate nearest neighbor searching fixed dimensions," *Journal on ACM*, vol. 45, no. 6, pp. 891–923, Nov. 1998.

[17] M. Muja and D. G. Lowe, "Fast approximate nearest neighbors with automatic algorithm configuration," in *International Conference on Computer Vision Theory and Application (VISSAPP)*, 2009, pp. 331–340.

[18] M. Fallon, H. Johannsson, M. Kaess, and J. Leonard, "The MIT stata center dataset," *International Journal of Robotics Research (IJRR)*, vol. 32, no. 14, pp. 1695–1699, December 2013.