

BoxGraph: Semantic Place Recognition and Pose Estimation from 3D LiDAR

Georgi Pramatarov, Daniele De Martini*, Matthew Gadd*, and Paul Newman
 Mobile Robotics Group (MRG), University of Oxford, **Equal contribution*
 {georgi,daniele,mattgadd,pnewman}@robots.ox.ac.uk

Abstract—This paper is about extremely robust and lightweight localisation using LiDAR point clouds based on instance segmentation and graph matching. We model 3D point clouds as fully-connected graphs of semantically identified components where each vertex corresponds to an object instance and encodes its shape. Optimal vertex association across graphs allows for full 6-Degree-of-Freedom (DoF) pose estimation and place recognition by measuring similarity. This representation is very concise, condensing the size of maps by a factor of 25 against the state-of-the-art, requiring only 3 kB to represent a 1.4 MB laser scan. We verify the efficacy of our system on the SemanticKITTI dataset, where we achieve a new state-of-the-art in place recognition, with an average of 88.4% recall at 100% precision where the next closest competitor follows with 64.9%. We also show accurate metric pose estimation performance – estimating 6-DoF pose with median errors of 10 cm and 0.33 deg.

I. INTRODUCTION

The ability to localise in an already visited location is crucial for autonomous operations. A commonly used approach is to build a map with the same sensor used during operations. Still, appearance change – due to change of viewpoint, occlusions caused by dynamic objects, or varying weather conditions – strongly affects the performances of such systems. LiDAR sensors are a very popular choice for autonomous vehicles due to their superior robustness to

these issues. Indeed, the fact that they actively sense the environment makes them less affected by appearance change and lighting conditions, and the geometric nature of their readings contains rich information about their surroundings.

Classically, LiDAR-based loop-closure detection methods rely on either local keypoints or global features. Unfortunately, the first emphasise local details – and result in huge maps – while the second discard many of the useful properties contained in the geometric distribution of the point cloud, often making metric estimation infeasible from the same data structures. Recent semantic-driven loop-closure detection define object-wise features, yielding *reliable and condensed* maps that contain both high-level information about the objects in the scene and their spatial relationships. A new vision-based approach [1] models input images of an indoor scene as fully-connected graphs of objects. These objects are detected by a network and described by the dimensions of their enclosing bounding boxes. Graph-matching and similarity measurement then amounts to finding optimal correspondences between objects.

Motivated by the success of this approach, we propose a global localisation method for 3D LiDAR point clouds that represents scenes as fully-connected semantic graphs (Fig. 1). Information on objects’ type and shape is

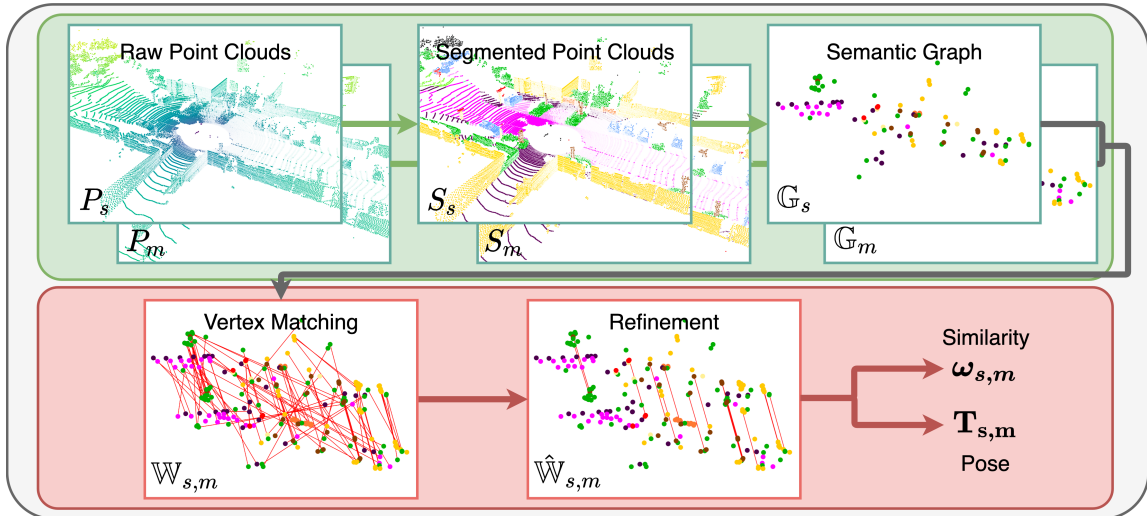


Fig. 1: System diagram of our method – raw point clouds, P_s and P_m , are labelled by a learned segmentation model and then clustered based on class to obtain compact complete graphs of semantic vertex instances, G_s and G_m . The instances are then matched according to their shape similarity, and the resulting correspondences $\mathbb{W}_{s,m}$ are refined to produce the final similarity score and pose estimate between the two scenes.

stored in vertices while their relative positions in the environment are stored in edges. Therefore, we capture both local and global geometrical features and higher-level semantic properties of a scene, resulting in powerful expressiveness. The place recognition task follows as an optimal assignment problem between the vertices of two graphs. In contrast to [1], we advance this method, showing that it can yield a 6-Degree-of-Freedom (DoF) metric pose estimate.

Our principal contributions are (1) the specialisation of a camera-based semantic graph-based place recognition system to 3D point clouds (using semantic blobs in 3D instead of 2D bounding boxes) and (2) extending this to allow for precise pose estimation on top of place recognition. We validate the approach on the KITTI odometry dataset and achieve against contemporary methods a *new state-of-the-art* in terms of place recognition (as compared to all approaches) and compactness (as compared to methods capable of both place recognition and metric pose estimation). We also achieve precise semantic-based pose estimation, where our performance is competitive with geometry-based approaches that operate on dense point clouds. Indeed, our evidence suggests that if given increasingly good segmentation predictions (which are not the focus of this paper), our method may challenge geometry-based methods in pose-estimation precision.

II. RELATED WORK

Traditional localisation techniques on 3D point clouds typically exploit low-level geometrical information. Approaches based on matching local keypoints [2], [3], [4] can yield precise pose estimates but rely on storing dense data and do not translate to place recognition. Matching of geometric primitives such as lines can be susceptible to occlusions [5]. Instead, global descriptor-based approaches abstract coarse information, becoming more susceptible to scene ambiguity, and are rarely coupled with accurate registration capabilities. For example, M2DP [6] projects a point cloud to several 2D planes and extracts density signatures for each plane, which could also be vulnerable to occlusions. Scan Context-based methods [7], [8] reduce a 3D point cloud to a flat polar grid, which could be affected by translational offsets. LiDAR-Iris [9] produces binary signature images from bird’s-eye-view projections of point clouds, also encoding the vertical distribution of points, which can suffer under roll and pitch variation. LocNet [10] proposes a histogram-based descriptor, contrastively tuned given similar scenes. However, this is still a much denser representation of the environment than our descriptor and only estimates a 3-DoF pose.

The advances of point-cloud learning methods such as PointNet [11] and PointNet++ [12] have shown their applicability in the domain of localisation: PointNetVLAD [13] and LPD-Net [14] combine learning approaches with the NetVLAD descriptor [15] to obtain a global scene representation; however, they discard information about the distribution of objects, and so they are not directly applicable to pose estimation. Segment-based methods such as SegMatch [16] and SegMap [17] partition point clouds into higher-level segments and, similarly to our approach, extract features for

each segment and match them across scenes. In contrast, however, they do not consider the relative positions between segments in each scan. Locus [18], also a segment-based method, extracts a global descriptor from topological and temporal segment features, which depends on sequential data aggregation, thus requiring motion before initialisation.

To improve the representational power of 3D descriptors, various methods employ semantics as additional inputs. Semantic Scan Context [19], for instance, extends [7] by encoding semantic information into the polar grid, demonstrating promising results; still, it only produces a rough 3-DoF pose estimate. OverlapNet [20] combines semantic, geometric, and intensity information and uses a Convolutional Neural Network (CNN) to estimate the relative yaw and overlap between laser scans. Instead, we explicitly extract information about the instances in a scene to create a higher-level – and compact – semantic and geometric representation.

SemSegMap [21], a semantic extension to [17], projects RGB and semantic information from an input camera feed to point clouds before building segments that show increased stability. They employ a contrastive learning scheme for segment feature extraction by self-supervising with segments from a pre-built map; however, this again relies on the aggregation of point clouds. PSE-Match [22] partitions point clouds per semantic class, extracts viewpoint-invariant spherical representations, and combines them into global place descriptors. They incorporate inter-class divergence learning at semantic instead of instance level, thus disregarding underlying relationships between objects of the same class.

Recently, graph-based approaches have shown potential in the global localisation domain when combined with semantic information. Both [23] and [24], similarly to us, create complete graphs of clustered instances from semantically-segmented point clouds. In [23], a neural network architecture is applied to extract vertex features based on graph connectivity and then predict the similarity of two graphs for place recognition. Thus, the system requires paired scans for supervision, while our approach can use pure geometric vertex features without training. GOSMatch [24], instead, computes histogram-based descriptors from the distances between the objects in the scene – a global graph descriptor for loop closure detection and vertex descriptors for initial pose estimation and verification. These approaches, however, consider only the spatial arrangement of object instances – via the graph edges. This motivates us to model both the relative distances and the individual objects’ shapes. We show that even a simple object shape descriptor (bounding box) can produce significant results in global localisation.

III. METHOD

Our approach to semantic localisation in the LiDAR domain is depicted in Fig. 1. We frame our system as a topometric localisation problem, where a point cloud

$$P_s : \{p_1 \dots p_n \mid p_i \in \mathbb{R}^3\} \quad (1)$$

from a LiDAR sensor stream \mathbb{P} is matched against point clouds P_m belonging to a pre-built map \mathbb{M} .




Fig. 2: Descriptor-matching amounts to finding the optimal assignment between same-class vertices based on shape similarity $\sigma_v(v_i)$. RANSAC removes outliers (e.g. v_4 to v_7 above) and results in a relative transformation $T_{s,m}$. The similarity score takes into account the similarity between matching vertices and edges of the refined graphs \hat{G}_s and \hat{G}_m : $\omega_{s,m} = \sigma_e(e_{12}, e_{56}) + \sigma_v(v_1, v_5) + \sigma_v(v_2, v_6)$.

We approach this problem by exploiting the semantic information in P_s and P_m and represent them as fully-connected graphs, G_s and G_m , to model objects – with their type and shape – and their relative positions. Each object is stored as a vertex encoding type, shape, and spatial location.

We then solve for localisation by finding the optimal association between the vertices of two such graphs, maximising geometric similarity. This produces both a 6-DoF pose estimate $T_{s,m} = [R_{s,m} | t_{s,m}]$ – where $R_{s,m}$ and $t_{s,m}$ are respectively the rotation matrix and the translation vector – and a final similarity score $\omega_{s,m}$ useful for place recognition.

A. Semantic Graph Extraction

To build a meaningful graph representation of a scene for localisation, distinctive semantic information needs to be extracted. This means a natural preference for stable objects – e.g. buildings and vegetation – over ephemeral and dynamic ones – such as vehicles and pedestrians – which we exclude. In contrast to [1] who use indoor object detection, we use semantically segmented point clouds as input to our system, since outdoor object detection and instance segmentation traditionally only focuses on dynamic objects.

Let \mathbb{L} be the set of the semantic classes of interest and define $\lambda : \mathbb{R}^3 \rightarrow \mathbb{N}$ to be a label function that can assign a semantic label $l \in \mathbb{L}$ to every point $p \in P$, resulting in a semantically enhanced point cloud

$$S = \{s_i \mid s_i = \{p_i, \lambda(p_i)\} \forall p_i \in P\} \quad (2)$$

where each $s_i \in S$ contains the semantic label l_i alongside the three-dimensional position of the corresponding point p_i .

We apply the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm [25] to cluster points with the same label based on their Euclidean distance. This results in a set of clusters of points, each constituting a point

cloud “local” to a region of the scan that shares the same semantic class

$$\begin{aligned} \mathbb{O} &= \{O_1, \dots, O_M \mid O_k \subset P, \\ &\quad l_k = \lambda(p_i) = \lambda(p_j) \forall p_i, p_j \in O_k\} \end{aligned} \quad (3)$$

Each cluster in \mathbb{O} is then used as a vertex to produce a fully-connected, semantic graph descriptor $G = \langle V, E \rangle$. A vertex $v_i = \langle c_i, l_i, f_i \rangle \in V$ is described by the geometrical centroid $c_i \in \mathbb{R}^3$ and the semantic class l_i of the corresponding cluster O_i , and a feature vector $f_i \in \mathbb{R}^D$ extracted from the cluster’s points and providing a compact representation of an object’s shape, as discussed further in Sec. III-B. An edge $e_{ij} = \langle v_i, v_j \rangle \in E$, instead, is described by the two centroids c_i and c_j of the vertices v_i and v_j respectively. In this view, the vertices model an approximation of the objects that compose the scene, whereas the edges model their relative position.

B. Vertex Feature Extraction and Vertex Similarity

To compare scenes and their associated semantic graphs, we need to extract objects’ shapes and measure their similarity. We use the feature vector of two vertices to compute the similarity between them, defined as $\sigma_v : V \times V \rightarrow \mathbb{R}$.

In contrast to [1], we start with 3D semantic blobs (not 2D detected image patches) and extract descriptors of the properties of scene entities. Bounding boxes provide a rough yet compact approximation of a cluster’s shape: intuitively, graphs of nearby scans result in similar clusters. In addition, the heading of a vehicle is typically aligned with the road in urban settings, so it is reasonable to assume that clusters’ bounding boxes will remain consistent across revisits. For this reason, we calculate the bounding box of each cluster and use its dimensions as the feature vector, i.e. $f_i = (h_i, w_i, d_i) \in \mathbb{R}^3$, where h_i , w_i , and d_i are the elongation of the cluster in the z , y and x directions respectively, measured along the LiDAR frame of reference mounted on the vehicle. If we let

$$\delta(a, b) = \frac{|a - b|}{\max(a, b)} \quad (4)$$

then a vertex similarity function σ_v can be defined as:

$$\sigma_v(v_i, v_j) = \begin{cases} \exp\left(-\frac{\delta(h_i, h_j) + \delta(w_i, w_j) + \delta(d_i, d_j)}{3}\right) & \text{if } l_i = l_j \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

Here δ computes the fractional variation of each dimension of the bounding boxes, and the average of such a discrepancy is used as the shape similarity metric.

C. Graph Matching

Let $G_s = \langle V_s, E_s \rangle$ and $G_m = \langle V_m, E_m \rangle$ be the graph representation of the point clouds $P_s \in \mathbb{P}$ and $P_m \in \mathbb{M}$. To calculate the similarity of the two semantic graphs, we need to find the optimal correspondence between their vertices. Let us assume for simplicity that $|V_s| \leq |V_m|$ without loss of generality. We seek to find a set of matches $\mathbb{W}_{s,m} = \{(v_i, v_j) \in V_s \times V_m\}$, where each v_i is matched to one and only one v_j , such that $\sum_{\mathbb{W}_{s,m}} \sigma_v(v_i, v_j)$ is

maximal. This is a classical instance of a linear assignment problem, which can be solved in polynomial time using the Hungarian algorithm, similarly to [1]. Moreover, as in [1], we improve the computational cost of this procedure by matching only the vertices of the same class, i.e. with non-zero similarity (Fig. 2).

D. Pose Estimation and Refinement

After obtaining the optimal vertex correspondences $\mathbb{W}_{s,m}$ between two graphs, in contrast with [1], we apply the Singular Value Decomposition (SVD) technique on vertex centroids to estimate a 6-DoF pose. As discussed above, however, the matching procedure associates each vertex of the smaller graph with one from the larger graph. Since this could lead to incorrect matches reducing the quality of the resulting pose estimate, we employ the RANdom SAmple Consensus (RANSAC) [26] algorithm to discard the outliers: we apply SVD iteratively on a subset of the matches to find the most inliers, subject to a tolerance threshold. The resulting transform represents the final pose estimate $T_{s,m} \in SE(3)$. In addition, the inliers form a refined set of matches $\hat{\mathbb{W}}_{s,m} \subseteq \mathbb{W}_{s,m}$, which we use next to calculate the graph similarity score (Fig. 2).

E. Similarity Score

Let $\omega_{s,m}$ be the similarity score between \mathbb{G}_s and \mathbb{G}_m . It is designed to take into account both the vertex similarities from the refined set – which include semantics and shape information – and the similarities of edges with matching endpoints – which contain the geometrical relationships between the objects; Fig. 2 shows a stylised example. In particular, the inlying set of matches $\hat{\mathbb{W}}_{s,m}$ is used to produce two fully-connected sub-graphs $\hat{\mathbb{G}}_s = \langle \hat{\mathbb{V}}_s, \hat{\mathbb{E}}_s \rangle \subseteq \mathbb{G}_s$ and $\hat{\mathbb{G}}_m = \langle \hat{\mathbb{V}}_m, \hat{\mathbb{E}}_m \rangle \subseteq \mathbb{G}_m$. As discussed above, each $e_{ij} = \langle v_i, v_j \rangle \in \hat{\mathbb{E}}_s$ and $e_{lk} = \langle v_l, v_k \rangle \in \hat{\mathbb{E}}_m$ are described by the centroids of their corresponding vertices. Similarly to [1], we define the edge similarity to be

$$\sigma_e(e_{il}, e_{jk}) = \exp(-\delta(\|c_i - c_l\|_2, \|c_j - c_k\|_2)) \quad (6)$$

where $e_{il} \in \hat{\mathbb{E}}_s$, $e_{jk} \in \hat{\mathbb{E}}_m$ and $\|\cdot\|_2$ is the l_2 -norm of a vector in \mathbb{R}^N . The final graph similarity score is then the sum of edge similarities of edges with matching vertices and the vertex similarities in the refined set:

$$\omega_{s,m} = \sum_{\hat{\mathbb{W}}_{s,m}} \sum_{\hat{\mathbb{W}}_{s,m}} \sigma_e(e_{il}, e_{jk}) + \sum_{\hat{\mathbb{W}}_{s,m}} \sigma_v(v_i, v_j) \quad (7)$$

$$\forall v_i, e_{il} \in \hat{\mathbb{G}}_s \text{ and } \forall v_j, e_{jk} \in \hat{\mathbb{G}}_m$$

We consider graphs to correspond to the same place if the similarity score between them is above some threshold $\tau \in \mathbb{R}$.

IV. EXPERIMENTAL RESULTS

We evaluate the proposed approach on the KITTI odometry dataset [27], which includes 11 sequences containing laser data from a Velodyne HDL-64 sensor and ground truth poses. For our experiments, we use the six sequences

Method	MPE	Descriptor shape		Size (kB)	
		Global	Vertex	Avg	Max
PointNetVLAD [13]	\times	256	-	1	1
M2DP [6]	\times	192	-	0.75	0.75
IRIS [9]	\times	80 × 360bit	-	3.52	3.52
ScanContext [7]	(✓)	20 × 60	-	4.69	4.69
SSC [19]	(✓)	50 × 360	-	70.3	70.3
OverlapNet [20]	(✓)	360 × 128	-	180	180
SGPR [23]	\times	100 × 4	-	1.56	1.56
GOSMatch [24]	(✓)	60 × 6	60 × 3	75.2	168.75
BG (Ours)	✓	-	7	2.87	6.51

TABLE I: Analysis of memory consumption of global localisation systems for storing a single scan using different place recognition methods. The baselines are categorised by whether they perform metric pose estimation (MPE), where (✓) denotes only a coarse estimate. Systems with \times and (✓) would also store the raw point cloud (not added in the table, 1.41 MB on average) and rely on other registration techniques to produce refined poses. Average and maximum number of vertices in a graph is taken to be 105 and 238.

that feature loop closures – 00, 02 and 05 to 08. In particular, Sequence 08 has reverse loops for which our system should yield consistent bounding boxes under the axis-aligned assumption.

We study our system’s performance with ground-truth semantic annotations from SemanticKITTI [28] and predictions from trained networks. In line with related work [19], [23], [24], we train a RangeNet++ [29] LiDAR segmentation model and provide an ablation analysis using predictions from a state-of-the-art network, Cylinder3D [30]. Indeed, while the latter achieves a mean Intersection over Union (mIoU) of 67.8 on SemanticKITTI, the former only reaches 52.2 and thus is key to assess the robustness of our approach to prediction noise and allows for a fair comparison with contemporary methods.

A. Map Compactness

As described in Sec. III, of 28 classes provided, we make use only of distinctive static objects (sidewalk, building, fence, vegetation, trunk, pole, and traffic-sign). For example, we leave out the road class as it is empirically less informative and frequently occluded. Semantic graphs are then generated for each scan and stored in map databases.

Throughout all the sequences, the average and maximum number of vertices in a graph are 105 and 238, respectively, with a total dimension of vertex descriptors of 7 – the centroid, the semantic label and the bounding box features. Tab. I includes the resulting dimensions of the graphs and compares the size of our representation with other methods.

As baselines, we select state-of-the-art place recognition methods based on geometric global descriptors such as PointNetVLAD [13], M2DP [6], LiDAR-Iris (IRIS) [9], and Scan Context [7], semantic global descriptors such as Semantic Scan Context (SSC) [19] and OverlapNet [20], as well as semantic graph variants (the family to which our method belongs), such as SGPR [23] and GOSMatch [24].

TABLE II: Aggregates of precision-recall performance across six SemanticKITTI sequences [28]. If based on semantics, methods from the upper group use predicted segmentation (e.g. SSC-RN uses RangeNet++ [29]). The lower group represent an upper limit on performance by using ground-truth segmentations. Our method is only outperformed in Sequence 08, which we prove in Tab. V is due to poor segmentation rather than the underlying localisation method – this is also clear from the best upper bound on performance demonstrated by BoxGraph (BG-SK).

Sequence →	00			02			05			06			07			08		
Method	F_1	$F_{0.5}$	F_2	F_1	$F_{0.5}$	F_2	F_1	$F_{0.5}$	F_2	F_1	$F_{0.5}$	F_2	F_1	$F_{0.5}$	F_2	F_1	$F_{0.5}$	F_2
PointNetVLAD [13]	.78	.82	.61	.73	.80	.56	.54	.63	.42	.85	.87	.68	.63	.72	.50	.04	.03	.05
M2DP [6]	.71	.78	.55	.72	.78	.56	.60	.71	.47	.79	.85	.59	.56	.69	.40	.07	.13	.07
IRIS [9]	.67	.70	.57	.76	.82	.60	.77	.83	.58	.91	.90	.74	.63	.69	.50	.48	.59	.38
ScanContext [7]	.75	.78	.62	.78	.82	.63	.89	.90	.72	.97	.94	.78	.66	.66	.54	.61	.65	.52
ISC [8]	.66	.73	.57	.71	.78	.59	.77	.83	.60	.84	.88	.65	.64	.69	.49	.41	.49	.35
SSC-RN [19]	.94	.92	.76	.89	.89	.72	.94	.93	.75	.99	.95	.79	.87	.89	.71	.88	.89	.71
OverlapNet [20]	.87	.87	.71	.83	.86	.64	.92	.91	.74	.93	.91	.75	.82	.80	.69	.37	.43	.32
SGPR [23]	.82	.81	.69	.75	.73	.65	.75	.78	.61	.65	.65	.59	.87	.88	.69	.75	.71	.65
<i>BG-RN (Ours)</i>	.99	.96	.79	.97	.94	.77	.96	.94	.76	1.00	.96	.80	1.00	.96	.80	.79	.85	.59
SSC-SK [19]	.95	.93	.77	.89	.88	.73	.95	.93	.76	.98	.95	.79	.88	.89	.72	.94	.93	.74
<i>BG-SK (Ours)</i>	1.00	.96	.80	.96	.94	.77	.97	.95	.77	1.00	.96	.80	1.00	.96	.80	.96	.94	.76

The authors of GOSMatch do not provide details on the exact Euclidean clustering procedure used to obtain their graph descriptor’s vertices, so we assume the resulting number of vertices is similar to ours. Some of these methods do not yield poses, while others only produce a coarse estimate which would require further refinement that relies on storing denser point cloud representations. We observe that the size of our graph representation is comparable to that of global-descriptor-based methods; yet, it also provides an accurate 6-DoF pose estimate (see Sec. IV-C). Compared to GOSMatch, which also estimates relative position between scans, we see a notable improvement in memory footprint of more than 25 times. With respect to the raw laser data, we reduce an average of 123 371 3D points per scan to only 105×7 , thus reducing memory consumption from 1446 kB to 2.87 kB.

B. Place Recognition

To assess the place recognition performance of our system, we estimate the similarity score $\omega_{s,m}$ (see Sec. III-E) between pairs of graphs from each sequence, in turn. We follow the same setup as in SSC [19] and SGPR [23], in

which a pair of scans are considered to be a true-positive place recognition if their relative distance is within 3 m, and a negative if the distance is outside 20 m. In particular, we use the benchmarks and evaluation pairs provided by SSC [19] which, for each sequence s , include all N_s positive pairs of scans more than 50 frames apart, and randomly sampled $100 \cdot N_s$ negative pairs. We benchmark our approach against PointNetVLAD [13], M2DP [6], LiDAR-Iris [9], Scan Context [7], Intensity Scan Context (ISC) [8], Semantic Scan Context (SSC) [19], OverlapNet [20] and SGPR [23]. We compare with SSC both using ground-truth labels from SemanticKITTI (SSC-SK) and predictions from RangeNet++ (SSC-RN). We show the precision-recall curves and present aggregate metrics, including maximum F-scores, given by

$$F_1 = 2 \frac{PR}{P + R}$$

$$F_\beta = (1 + \beta^2) \frac{PR}{\beta^2 P + R}$$

TABLE III: Aggregates of precision-recall performance across six SemanticKITTI sequences, please refer to Tab. II for instructions on how to read these results. In this, R_1 refers to recall at 100 % precision.

Sequence →	00			02			05			06			07			08		
Method	R_1	AP	EP	R_1	AP	EP	R_1	AP	EP	R_1	AP	EP	R_1	AP	EP	R_1	AP	EP
PointNetVLAD [13]	.28	.81	.64	.38	.73	.69	.07	.54	.54	.53	.89	.77	.18	.67	.59	.50	.01	.25
M2DP [6]	.23	.73	.62	.21	.71	.60	.22	.62	.61	.36	.77	.68	.17	.49	.59	.50	.04	.25
IRIS [9]	.25	.74	.63	.33	.79	.67	.49	.80	.75	.58	.97	.79	.30	.66	.65	.12	.48	.56
ScanContext [7]	.22	.83	.61	.26	.82	.63	.59	.95	.80	.85	.99	.92	.11	.66	.55	.14	.64	.57
ISC [8]	.25	.74	.63	.22	.76	.61	.45	.83	.73	.63	.88	.82	.28	.66	.64	.09	.40	.54
SSC-RN [19]	.65	.99	.83	.49	.95	.74	.80	.98	.90	.95	.99	.97	.54	.95	.77	.46	.95	.73
OverlapNet [20]	.11	.93	.56	.28	.84	.64	.59	.96	.80	.49	.98	.74	.17	.89	.59	.00	.29	.17
SGPR [23]	.00	.88	.50	.50	.74	.75	.06	.80	.53	.05	.68	.48	.44	.91	.72	.04	.78	.52
<i>BG-RN (Ours)</i>	.98	1.00	.99	.86	.98	.93	.91	.96	.95	.99	1.00	.99	.99	1.00	1.00	.58	.76	.79
SSC-SK [19]	.70	.99	.85	.50	.96	.75	.81	.99	.90	.94	1.00	.97	.61	.96	.80	.86	.97	.93
<i>BG-SK (Ours)</i>	.98	1.00	.99	.78	.98	.89	.94	.97	.97	.99	1.00	1.00	1.00	1.00	1.00	.87	.96	.93

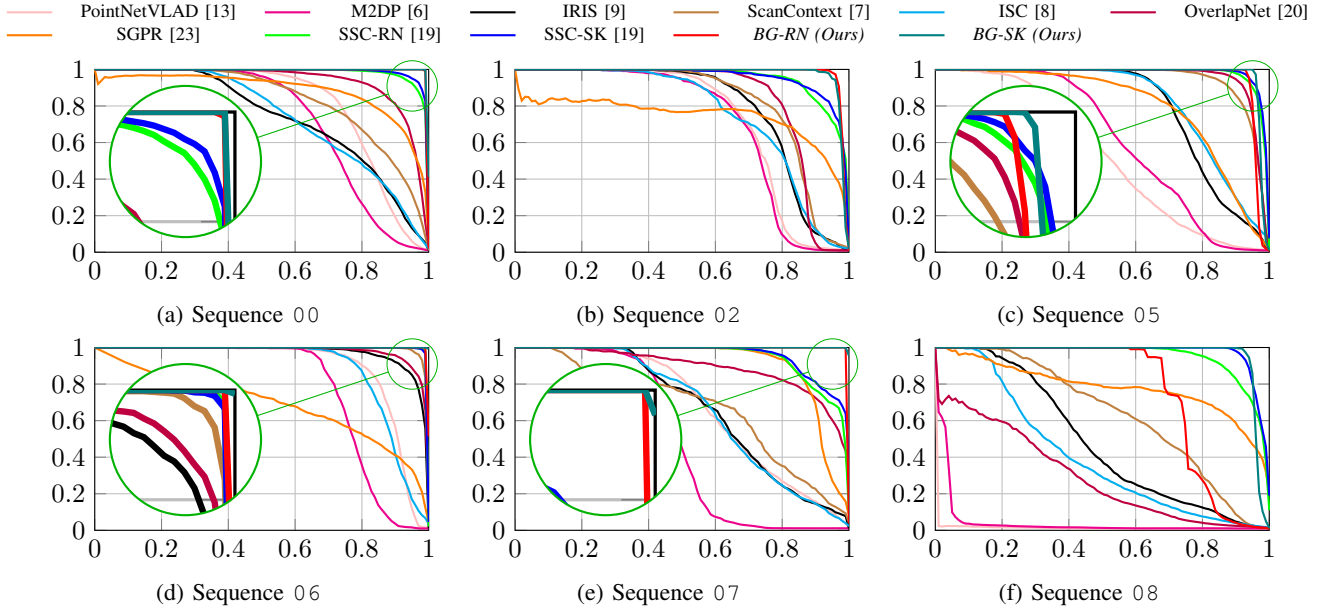


Fig. 3: Precision-recall curves across six Semantic KITTI Sequences [28]. These curves are aggregated in Tabs. II and III. Cluttered results are presented on a zoomed in scope, for better readability.

as well as recall at 100% precision (R_1), Average Precision (AP) and Extended Precision [31] (EP),

$$AP = \sum_n (R_n - R_{n-1})P_n, \quad EP = \frac{1}{2}(R_1 + P_0)$$

where P_0 is the precision at 0% recall.

Fig. 3 and Tabs. II and III show that we outperform all other methods in every aggregate and over every sequence. For example, in Sequence 00, we obtain as much as 0.98 R_1 . To be clear, this means that the system may be tuned to make *no mistakes* across 98% of a large-scale urban autonomous operation, where we only allow it to assume that the single most similar scan is a *bona fide* match. The worst-performing sequence on RangeNet++ predictions is Sequence 08, which contains entirely reverse loop closures. The recall here is 0.58. Using ground truth, however, demonstrates comparable results with the other sequences – 0.87 – suggesting that the lesser performance is primarily due to incorrect predictions by the network (see Sec. IV-D). Contrast this with M2DP and OverlapNet, which are degraded under revisits in the opposite direction.

We obtain a max F_1 score of 0.951 and R_1 of 0.884 for *BG-RN (Ours)*, averaged across all sequences. This is in contrast with 0.766 and 0.182, respectively, for SGPR, another semantic-graph-based approach. Low recall means that SGPR would be prone to false positive loop closures. Our closest competitor is SSC-RN with average maximum F_1 score of 0.918 and R_1 of 0.649. Importantly, we outperform SSC-SK (which uses ground truth segmentations), even when we do not use ground truth segmentations (BG-RN).

C. Pose Estimation

Similarly to Sec. IV-B, we evaluate the metric localisation capabilities of our system by estimating the pose $T_{s,m}$ (see Sec. III-D) between pairs of graphs from each sequence database, in turn. Here we consider all pairs of distinct scans within ground truth distance of 3 m. We found the ground-truth GPS poses of the KITTI odometry dataset to be noisy, particularly for Sequence 00, so we refine them using the Iterative Closest Point (ICP) algorithm. We compare our performance with two global registration approaches based on matching dense local features with RANSAC – using the




Fig. 4: Quartile (q_1 , q_2 , and q_3) and full-range (min / max) RTE statistics for metric pose estimation across six SemanticKITTI sequences [28]. Note that the results for SSC-SK (red) are generated using *ground-truth* semantic segmentations, while our method (*BG-RN*, blue) uses predictions. We nevertheless consistently outperform SSC as well as FPFH (orange), except in Sequence 08. Tab. V shows this is due to the poor quality of input predictions. It is also important to note that with ideal segmentations (*BG-SK*), our method is capable of outperforming even FCGF [4] (black), a state-of-the-art learned approach.

TABLE IV: RTE (top, values in m) and RRE (bottom, in deg) pose estimation results, showing median (q_2) and lower and upper quartiles (q_1 and q_3) per SemanticKITTI sequence. Bold shows best performing, while underlined shows second-best performing. *BG-RN (Ours)* from the upper groups uses predicted segmentation (FPFH and FCGF are not based on semantics). The lower groups represent an upper limit on performance by using ground-truth segmentations. Even though SSC uses ground truth (i.e. SSC-SK), our system lacking ground truth (*BG-RN*), consistently outperforms it by a large margin. It is important to note that the upper limit on performance for our system *BG-SK* in fact outperforms the learned FCGF.

Sequence →	00			02			05			06			07			08		
Method	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3
FPFH [2]	.13	.09	.19	.14	.09	.22	.12	.08	.19	.17	.11	.26	.11	.07	.17	<u>.16</u>	.10	<u>.25</u>
FCGF [4]	.08	.05	.11	.09	.06	.15	.07	.05	.10	.06	.05	.09	<u>.08</u>	.05	.11	.08	.05	.12
<i>BG-RN (Ours)</i>	<u>.10</u>	<u>.06</u>	<u>.17</u>	<u>.11</u>	<u>.07</u>	<u>.19</u>	<u>.09</u>	.05	<u>.15</u>	<u>.11</u>	<u>.07</u>	<u>.19</u>	.07	<u>.05</u>	<u>.12</u>	<u>.16</u>	<u>.09</u>	.48
SSC-SK [19]	.17	.09	.34	.21	.10	.47	.15	.07	.33	.19	.11	.36	.17	.07	.44	.18	.09	.36
<i>BG-SK (Ours)</i>	.08	.05	.12	.08	.05	.13	.07	.04	.10	.06	.04	.09	.06	.04	.10	.08	.05	.12

Sequence →	00			02			05			06			07			08		
Method	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3	q_2	q_1	q_3
FPFH [2]	<u>.35</u>	.22	<u>.56</u>	.38	.24	<u>.64</u>	.32	.19	.54	.39	.23	.64	.28	.16	<u>.47</u>	<u>.46</u>	.27	<u>.84</u>
FCGF [4]	.18	.11	.29	.21	.13	.38	.14	.08	.24	.11	.07	.16	.16	.09	.28	.17	.10	.30
<i>BG-RN (Ours)</i>	.36	<u>.19</u>	.68	<u>.34</u>	<u>.18</u>	.70	<u>.25</u>	<u>.10</u>	<u>.52</u>	<u>.26</u>	<u>.13</u>	<u>.51</u>	<u>.22</u>	<u>.08</u>	<u>.47</u>	.49	<u>.22</u>	1.35
SSC-SK [19]	.56	.32	.95	.46	.26	.84	.50	.21	.86	.40	.19	.73	.44	.14	.90	.58	.30	1.05
<i>BG-SK (Ours)</i>	.29	.16	.50	.26	.14	.49	.21	.10	.37	.16	.10	.24	.20	.08	.39	.27	.15	.50

classical Fast Point Feature Histograms (FPFH) [2], and the learned Fully Convolutional Geometric Features (FCGF) [4]. For FPFH and FCGF, we downsample input point clouds with a voxel size of 30 cm and we use FCGF’s pre-trained model on KITTI with 32 feature dimensions. For a fair comparison, we set the RANSAC iterations for all methods to 10 000. We also compare with SSC [19] using ground-truth labels as a representative of a fellow semantic-based loop-closure approach that also yields an initial pose, but it is worth noting this pose has only 3-DoF. We report the Relative Translation Error (RTE) and Relative Rotation Error (RRE) following the definitions from [4]:

$$RTE = \|t_{s,m} - t_{s,m}^*\|_2, \quad RRE = \cos^{-1}\left(\frac{\text{Tr}(R_{s,m}^T R_{s,m}^*) - 1}{2}\right)$$

where R^* and t^* are the ground-truth values for the estimated pose and $\text{Tr}(\cdot)$ is the trace operation.

Here, Fig. 4 shows per-sequence RTE error distributions as box-and-whisker plots. We can see that our method performs consistently across all sequences, except Sequence 08 when using predictions, which we deduce is due to prediction noise (see Sec. IV-D). SSC-SK performs worst, which suggests it would require further refinement when deployed on a live system. We can conclude that out of all the place recognition approaches in Sec. IV-B, we are the only one that also produces an accurate pose estimate. Tab. IV shows the median as well as lower and upper quartile of RTE and RRE. Aggregated over *all sequences*, we have median RTE of 13.4 cm for FPFH, 7.7 cm for FCGF, 17 cm for SSC-SK (with ground-truth segmentations), 10.4 cm for BG-RN (ours with *predicted* segmentations), 7.2 cm for BG-SK (ours with ground-truth segmentations). Similarly, median RRE is 0.361° for FPFH, 0.167° for FCGF, 0.513° for SSC-SK, 0.331° for BG-RN, 0.246° for BG-SK. Note that FPFH and

FCGF operate on downsampled scans with around 20 000 points (234.38 kB), which would still need to be stored in the map. As mentioned in Tab. IV, when using predictions, we are outperforming non-learned geometry-based methods and approaching the performance of learned variants. Indeed, it would seem that our method, if presented with ideal segmentations, is capable of better performance than FCGF, a descriptor trained explicitly to register LiDAR scans.

D. Effect of segmentation prediction quality

We investigate the effect of the performance of the prerequisite segmentation network on our localisation capability, where ground-truth segmentation labels provide an upper limit on performance. This is done both to understand the applicability of various trained networks and to confirm the efficacy of our underlying localisation formulation beyond any variance in segmentation network performance. We focus on Sequence 08, the *only* sequence with a considerable drop in performance when using predicted as opposed to ground-truth labels Figs. 3 and 4.

Tab. V shows that our method is improved in its place recognition performance by the use of Cylinder3D [30] – a better segmentation network – achieving for example a maximum F_1 score of 0.89 (originally 0.79). Correspondingly, our median RTE improves from 0.16 to 0.12.

This, alongside the evidence in Tabs. II and III in which *BG-SK* outperforms *SSC-SK across all aggregates and all sequences*, further suggests that our upper limit on performance is superior to our close competitor. Importantly, *BG-RN* is generally more robust even with poorer segmentation predictions for all other sequences in Tabs. II and III.

TABLE V: Effect on both place recognition and metric pose estimation performance of using better segmentation predictions from the input network (RN – RangeNet++ [29], CYL3D – Cylinder3D [30], SK – ground truth).

Method	Place recognition						Pose estimation	
	F_1	$F_{1.5}$	F_2	R_1	AP	EP	$RTE(m)$	$RRE(o)$
BG-RN	.79	.85	.59	.58	.76	.79	.16	.49
BG-CYL3D	.89	.91	.68	.72	.88	.86	.12	.40
BG-SK	.96	.94	.76	.87	.96	.93	.08	.27

V. CONCLUSION

We have presented a system for extremely lightweight and performant LiDAR localisation using a compact but discriminative representation derived from semantic segmentation of raw laser data. We adapt a semantic graph descriptor to high-fidelity 3D laser scans for both place recognition and rigid-body registration. We apply the representation to place recognition using 3D bounding boxes as shape embeddings of vertex entities and extend the method to pose estimation. We demonstrate state-of-the-art place recognition with competitive metric localisation performance on the KITTI odometry dataset while representing scans with a fractional memory requirement. Our method, if presented with ideal segmentations, is capable of better performance than systems trained explicitly to estimate precise pose between dense laser scans, and we expect that, as scene understanding theory and practice develop, the work presented in this paper will be exploited heavily in robust localisation with LiDAR.

ACKNOWLEDGEMENTS

Thanks to the Assuring Autonomy International Programme, a partnership between Lloyd’s Register Foundation and the University of York, and EPSRC Programme Grant “From Sensing to Collaboration” (EP/V000748/1).

REFERENCES

- [1] C. Qin, Y. Zhang, Y. Liu, and G. Lv, “Semantic loop closure detection based on graph matching in multi-objects scenes,” *Journal of Visual Communication and Image Representation*, 2021.
- [2] R. B. Rusu, N. Blodow, and M. Beetz, “Fast point feature histograms (fpfh) for 3d registration,” in *International Conference on Robotics and Automation*, 2009.
- [3] F. Kallasi, D. L. Rizzini, and S. Caselli, “Fast keypoint features from laser scanner for robot localization and mapping,” *Robotics and Automation Letters*, 2016.
- [4] C. Choy, J. Park, and V. Koltun, “Fully convolutional geometric features,” in *International Conference on Computer Vision*, 2019.
- [5] T. Suleymanov, M. Gadd, L. Kunze, and P. Newman, “Lidar lateral localisation despite challenging occlusion from traffic,” in *Position, Location and Navigation Symposium*, 2020, pp. 334–341.
- [6] L. He, X. Wang, and H. Zhang, “M2dp: A novel 3d point cloud descriptor and its application in loop closure detection,” in *International Conference on Intelligent Robots and Systems*, 2016.
- [7] G. Kim and A. Kim, “Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map,” in *International Conference on Intelligent Robots and Systems*, 2018.
- [8] H. Wang, C. Wang, and L. Xie, “Intensity scan context: Coding intensity and geometry relations for loop closure detection,” in *International Conference on Robotics and Automation*, 2020, pp. 2095–2101.
- [9] Y. Wang, Z. Sun, C.-Z. Xu, S. E. Sarma, J. Yang, and H. Kong, “Lidar iris for loop-closure detection,” in *International Conference on Intelligent Robots and Systems*, 2020, pp. 5769–5775.

- [10] H. Yin, Y. Wang, X. Ding, L. Tang, S. Huang, and R. Xiong, “3d lidar-based global localization using siamese neural network,” *Transactions on Intelligent Transportation Systems*, vol. 21, no. 4, 2020.
- [11] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” in *Conference on Computer Vision and Pattern Recognition*, 2017.
- [12] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “Pointnet++: Deep hierarchical feature learning on point sets in a metric space,” *Advances in neural information processing systems*, vol. 30, 2017.
- [13] M. A. Uy and G. H. Lee, “Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition,” in *Conference on Computer Vision and Pattern Recognition*, 2018.
- [14] Z. Liu, S. Zhou, C. Suo, P. Yin, W. Chen, H. Wang, H. Li, and Y. Liu, “Lpd-net: 3d point cloud learning for large-scale place recognition and environment analysis,” in *International Conference on Computer Vision*, 2019.
- [15] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, “Netvlad: Cnn architecture for weakly supervised place recognition,” in *Conference on Computer Vision and Pattern Recognition*, 2016.
- [16] R. Dubé, D. Dugas, E. Stumm, J. Nieto, R. Siegwart, and C. Cadena, “Segmatch: Segment based place recognition in 3d point clouds,” in *International Conference on Robotics and Automation*, 2017.
- [17] R. Dubé, A. Cramariuc, D. D. H. Sommer, M. Dymczyk, J. N. R. Siegwart, and C. Cadena, “SegMap: Segment-based mapping and localization using data-driven descriptors,” *The International Journal of Robotics Research*.
- [18] K. Vidanapathirana, P. Moghadam, B. Harwood, M. Zhao, S. Sridharan, and C. Fookes, “Locus: Lidar-based place recognition using spatiotemporal higher-order pooling,” in *International Conference on Robotics and Automation*, 2021, pp. 5075–5081.
- [19] L. Li, X. Kong, X. Zhao, T. Huang, W. Li, F. Wen, H. Zhang, and Y. Liu, “Ssc: Semantic scan context for large-scale place recognition,” in *International Conference on Intelligent Robots and Systems*, 2021.
- [20] X. Chen, T. Labe, A. Milioto, T. Röhling, J. Behley, and C. Stachniss, “OverlapNet: A Siamese Network for Computing LiDAR Scan Similarity with Applications to Loop Closing and Localization,” *Autonomous Robots*, vol. 46, pp. 61–81, 2021.
- [21] A. Cramariuc, F. Tschopp, N. Alatur, S. Benz, T. Falck, M. Brühlmeier, B. Hahn, J. Nieto, and R. Siegwart, “Sensesegmap–3d segment-based semantic localization,” in *International Conference on Intelligent Robots and Systems*, 2021, pp. 1183–1190.
- [22] P. Yin, L. Xu, Z. Feng, A. Egorov, and B. Li, “Pse-match: A viewpoint-free place recognition method with parallel semantic embedding,” *Transactions on Intelligent Transportation Systems*, 2021.
- [23] X. Kong, X. Yang, G. Zhai, X. Zhao, X. Zeng, M. Wang, Y. Liu, W. Li, and F. Wen, “Semantic graph based place recognition for 3d point clouds,” in *International Conference on Intelligent Robots and Systems*, 2020.
- [24] Y. Zhu, Y. Ma, L. Chen, C. Liu, M. Ye, and L. Li, “Gosmatch: Graph-of-semantics matching for detecting loop closures in 3d lidar data,” in *International Conference on Intelligent Robots and Systems*, 2020.
- [25] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *kdd*, vol. 96, no. 34, 1996, pp. 226–231.
- [26] M. A. Fischler and R. C. Bolles, “Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography,” *Commun. ACM*, vol. 24, pp. 381–395, 1981.
- [27] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *Conference on Computer Vision and Pattern Recognition*, 2012.
- [28] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, “Semantickitti: A dataset for semantic scene understanding of lidar sequences,” in *International Conference on Computer Vision*, 2019.
- [29] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, “Rangenet++: Fast and accurate lidar semantic segmentation,” in *International Conference on Intelligent Robots and Systems*, 2019.
- [30] H. Zhou, X. Zhu, X. Song, Y. Ma, Z. Wang, H. Li, and D. Lin, “Cylinder3d: An effective 3d framework for driving-scene lidar semantic segmentation,” *arXiv preprint arXiv:2008.01550*, 2020.
- [31] B. Ferrarini, M. Waheed, S. Waheed, S. Ehsan, M. J. Milford, and K. D. McDonald-Maier, “Exploring performance bounds of visual place recognition using extended precision,” *Robotics and Automation Letters*, vol. 5, no. 2, pp. 1688–1695, 2020.