Learning to See the Wood for the Trees:
Deep Laser Localization in Urban and Natural Environments on a CPU

Georgi Tinchev, Adrian Penate-Sanchez and Maurice Fallon

Abstract— Localization in challenging, natural environments such as forests or woodlands is an important capability for many applications from guiding a robot navigating along a forest trail to monitoring vegetation growth with handheld sensors. In this work we explore laser-based localization in both urban and natural environments, which is suitable for online applications. We propose a deep learning approach capable of learning meaningful descriptors directly from 3D point clouds by comparing triplets (anchor, positive and negative examples). The approach learns a feature space representation for a set of segmented point clouds that are matched between a current and previous observations. Our learning method is tailored towards loop closure detection resulting in a small model which can be deployed using only a CPU. The proposed learning method would allow the full pipeline to run on robots with limited computational payload such as drones, quadrupeds or UGVs.

I. INTRODUCTION

Localization is a fundamental task in robotic perception, a robot needs to know where it is to navigate in the environment and to make decisions. It has been heavily explored with computer vision, resulting in the visual teach-and-repeat concept demonstrating impressive results at large scales [1], [2]. These types of approaches typically assume a certain inherent structure in the scene, image features are dependant on repeatable camera viewpoint [3], and methods are usually tested in urban environments which guide the robot along the route in question [4]. While the aforementioned visual approaches have many promising characteristics, here we explore LIDAR due to its robustness to varying lighting conditions, changes in viewpoint, and trackline offsets. It is a precise and long range sensing modality.

The SegMatch system [5] proposed a modular segment-based approach for LIDAR teach-and-repeat, which could localize within a prior map while also retaining a degree of semantic meaning, as the segments matched corresponded to large physical objects such as cars and parts of a building. The approach was further improved with learned feature descriptors [6] which achieved greater accuracy. The approach uses a GPU to achieve real-time performance.

We are motivated to localize in natural scenarios like forests. In these environments many assumptions about the appearance or the geometry of landmarks are no longer valid - vegetation grows between seasons and there is no
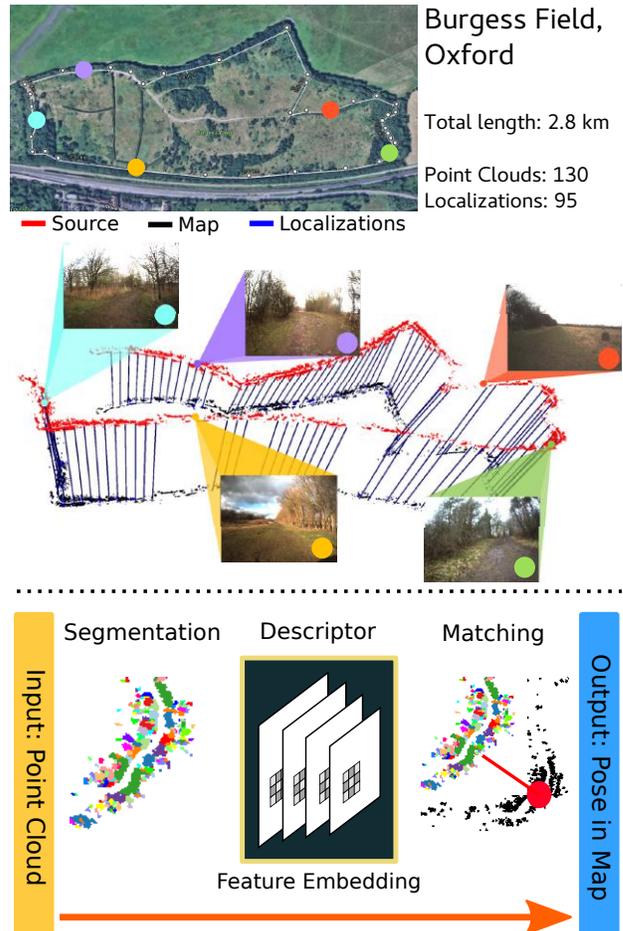
Fig. 1: **Top:** Depiction of the obtained results on large, unstructured environments. The proposed approach regularly proposes correct loop closures despite the challenging natural scene, using a deep learning architecture only on a CPU. **Bottom:** Overview of the used methodology. Segments are extracted from the input point cloud on which our neural network features are obtained and finally matched and the pose estimated.

planar structure. Our previous work showed promising results in both structured (urban) and unstructured (natural) environments [7]. While reliable localization was achieved, we relied on a hand-crafted set of features which limited the performance.

The main contribution of this work is a novel method for segment-based LIDAR localization. We aim to improve upon our previous work by learning segment representations in a way that inherently handles the variability of the given

environment. The proposed approach learns a descriptor space which efficiently represents the similarities between partial observations of the same segment which makes it robust to incomplete data. We use a neural network to learn this high dimensional feature space. The proposed approach utilizes the convolution operation proposed in [8] to learn an embedding space for both urban and natural scenarios directly from the raw point cloud data. The neural network has the following characteristics:

- Unordered point clouds as input: does not require a specific ordering of the point clouds in a segment. This makes our approach flexible by avoiding the computational cost of creating specific structures for the point cloud and aligning the inputs to a grid or voxelization.
- Feature space is capable of being generalized: experimental validation has proven that our proposed deep learning solution creates a feature space that can generalize, without the need to be retrained to a new sensor or a different environment.
- The network can estimate the quality of a match: a probability is computed and can be used when carrying out probabilistic geometric validation such as PROSAC [9], making our approach more efficient.

In the results section we demonstrate that our proposed method significantly outperforms other hand-engineered approaches, while also improving the computational speed in comparison to other deep learning approaches. In particular we achieve localization performance similar to SegMap [6], but do not require a GPU at runtime. Instead the method can be deployed online on the CPU of a mobile robot. This performance improvement comes in large part because our network is more specifically tailored to the task of localization. As presented in [6], the SegMap learning approach can be used to compress and reconstruct point clouds as well as extract and use semantic meaning from the segments to aid localization. We primarily focus on localization in natural scenes, where our approach is more capable of expressing the details of a point cloud [10], [11], [8].

II. RELATED WORK

Robot localization has been heavily explored using different sensors such as LIDAR [12], vision [1], GPS [13] or radar [14] with reliable approaches often combining different sensing modalities [12], [13], [14]. Our proposed work focuses on global LIDAR localization particularly applied to unstructured environments rather than incremental localization or odometry. Here we will briefly overview some methods performing LIDAR localization.

For self-driving vehicles many approaches have taken advantage of LIDAR reflectivity to achieve precise localization. These methods are commonly used in commercial approaches, for example, the approach of [15] uses a prior map of reflectivity and exploits road marks to reliably localize a vehicle in an urban environment. The authors formulate the world as a mixture of Gaussians (GMM) over 2D grid cells. The GMM represents the heights of points in each cell and the reflectance in a vigorous way that allows

the approach to be robust to weather alteration and road degradation.

When localization is carried out concurrently with mapping it is often treated as a research field of its own (Simultaneous Localization and Mapping, SLAM) where the key problem of loop-closure detection is often equivalent to global localization. In this context the robot's odometry provides a rough estimate of the position of the robot within the world which can be leveraged to avoid searching for loop closures over the entire map, thus keeping the computational cost low [16]. When the assumption of a local neighbourhood is lost most of these approaches require ICP alignments or similar costly alternatives [17].

Recent approaches aim to localize on a higher level than on explicit point-to-point basis. The concept of segment localization was presented in [18] but more specifically our work is directly motivated by the works of Dubé *et al* [5], [6]. We will describe the framework of segment-based localization in Section III. Features created from LIDAR data are less informative or unique than visual features, thus approaches later prune the candidate matches to reduce the percentage of outliers before performing a robust geometric validation [19]. One of the key advantages of using segments instead of keypoints is that more semantically meaningful entities are extracted, increasing the repeatability of the descriptors. This was demonstrated in [5] and was applied to urban environments. In [6] segment localization evolved into a segment-based SLAM that introduced a new segment feature that also encoded semantics and volumetric shape. This interesting line of work learns features suitable for localization of voxelized segments and the semantics of the map simultaneously. The approach was evaluated in realistic urban setting and in derelict buildings. Elbaz *et al* [20] used segments produced from a Random Sphere Cover Set - overlapping point cloud spheres where each point of the original point cloud could be part of more than one sphere. These spheres were then projected to 2D depth images and processed by a Deep Auto-Encoder. In our previous work [7] we proposed a larger hybrid feature descriptor which improved detection performance and achieved localizations in unstructured environments such as forests, while also being capable of handling the urban scenarios.

The main drawback of classical approaches that use LIDAR data is that they rely on handcrafted features which do not perform as well as learned features [21]. When trying to apply learning techniques to 3D data there are two prominent approaches - representing the data as a voxelized grid, or by using the more common representation of point clouds. Initially these approaches tended to use voxelized inputs as the tensor operators generalize directly to this representation and achieve promising results as shown in [22]. The main drawback of this approach is the amount of computation required to model the entirety of the environment. In [22] the authors used a 32x32x32 voxelized grid, resulting in a dimensionality of 32768, with 3D convolutions applied in strides. Because the input resolution is low (due to the computational complexity), these approaches are not capable
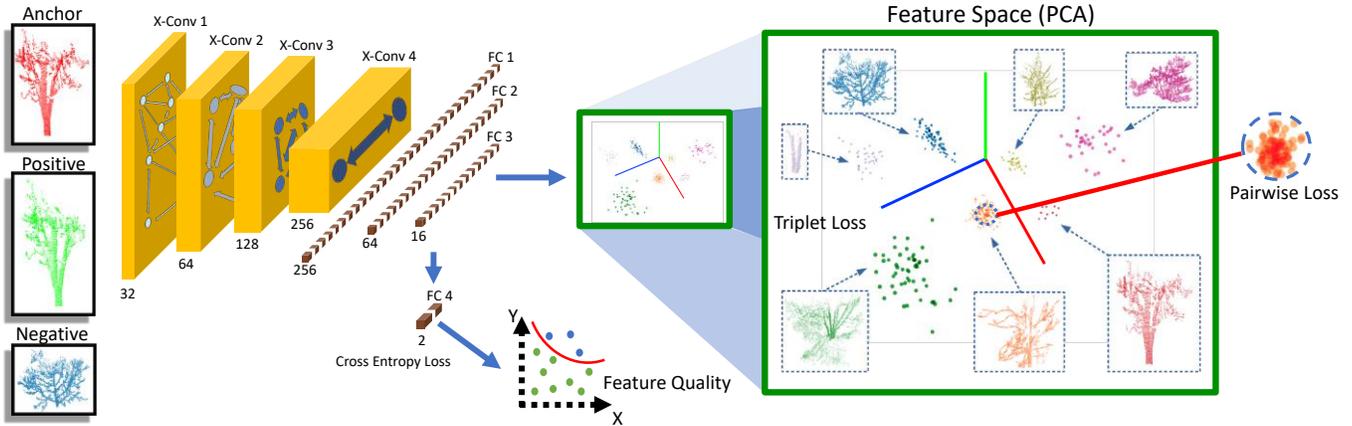
Fig. 2: Network architecture used for learning the feature embeddings (left) and a visualization of the embedding space after training (right). The network takes as input raw point data, it uses four x-convolutional layers ($\mathcal{X}$-conv) and three fully connected (FC) layers to estimate the feature space. A feature quality branch is also created by using a fully connected layer appended to the feature embedding. The networks learns to cluster similar objects (denoted in the same colour) close together while separating dissimilar objects. The final feature descriptor is trained using a combination of triplet and pairwise losses.

of handling detailed shapes such as the legs of a chair, for example, as shown in [10]. The issue of handling point clouds was addressed in [23], [11] by learning a symmetry function approximated by a multi-layer perceptron (MLP). The authors created a neural network layer that was invariant to point order. In contrast to the aforementioned volumetric approach, a raw processed point cloud containing 1024 points resulted in a dimensionality of only 3072. In addition, by using a MLP as the basis of their layer and by aggregation of spatial data, Li *et al* [8] managed to create a more descriptive layer to handle unordered point clouds.

III. METHODOLOGY

In this section we present Efficient Segment Matching (**ESM**). The goal of the algorithm is to estimate the position of a robot with respect to a prior map. The task is split into four modules - Segmentation, Description, Matching and Pose Estimation. This work retains the approach to pose estimation which was initially proposed by [5]. Our contributions in this paper are (1) learning a novel descriptor of segments together with (2) a learned measure of performance per descriptor which enables the pruning of the features before matching and (3) the use of a probabilistic robust pose estimation [9]. The Segmentation, Matching and part of the Pose Estimation have remained unaltered from our previous work so we will offer an overview of them and rather focus on the novel contributions of this work. An input point cloud is first filtered and segmented to produce a number of 3D point cloud objects that are matched between an a-priori map and the current live point cloud. All the segments are passed to a neural network to extract the descriptor of each segment and a measure of the quality of that descriptor. For each segment in the live point cloud the N-closest neighbors are found based on the Euclidean distance of the descriptor compared to the descriptors of the segments from the map. These matches are ordered based on the extracted feature

quality and passed to PROSAC [9] to produce a set of possible localizations. In this work the ordering is done by computing the joint probability of both features (the feature from the live point cloud and the one from the map). The order of the matches is important because the geometric consistency and PROSAC consider first the matches with higher likelihood. Each module of this pipeline is described in depth in our previous work [7]. In the following section we detail the architecture and the learning approach.

A. Architecture

Our network architecture is constructed based on the novel convolutional layers presented in [8]. Fig. 2 illustrates our model. The network learns directly from point cloud data, it takes as input a batch of raw point clouds. Each raw point cloud is uniformly downsampled, zero-centered, and scaled. The network used for our descriptor learning approach consists of four $\mathcal{X}$-conv operators [8] and three fully connected layers; dropout [24] of 0.5 is applied at the second fully connected layer. The weights at the last fully connected layer are used as the descriptor. The $\mathcal{X}$-conv operator convolves local regions, similar to convolutions in images by a CNN. For each point, its closest N neighbors are projected to a local coordinate frame and lifted to a higher dimensional space with a Multi-Layer Perceptron (MLP). The $\mathcal{X}$-conv operator learns a convolution based on the MLPs of the neighbouring points. In order for the top $\mathcal{X}$-conv operators to see a larger portion of the point cloud a dilation rate (D) is applied [25]. In this way the receptive field of the top layers is increased without an increase of the neighbouring point size (N) or the kernel size. We have selected a specific configuration of the dilation and neighbourhood in each of our layers that better represents the problem we are trying to solve. The first layer uses a neighbourhood of 8 and a dilation of 1 per $\mathcal{X}$-conv operator, the second a neighbourhood of 12 and dilation 3, the third

a neighbourhood of 16 and a dilation of 3 and the fourth a neighbourhood of 16 and a dilation of 4. The reason for this is that the first layer will look at a few points (neighbourhood 8) that are immediately close (dilation 1). The second will look into more points (neighbourhood 12) further out (dilation 3). This way the network creates a representation that aggregates the information slowly creating a hierarchical representation of the whole segment. This is done for each point so as to create a feature of the entire segment from the relative viewpoint of each point. The final descriptor is computed as an average of the features for all the points in a single point cloud. By doing this the feature fuses data from different viewpoints creating an expressive yet simple representation. In all our experiments we managed to achieve good performance by using an embedding size of 16 dimensions. The dimensionality of a descriptor has a strong impact on the performance of matching. This keeps the computational cost low even if the segmentation algorithm produces many more segments.

The proposed network architecture also includes a classification branch that estimates a measure of the quality of each descriptor. To be able to train this branch we need to train the feature network first. When training the classification branch the descriptor layers are not modified, the model optimizes a single fully connected layer of size two. This represents a logistic regression that classifies whether a feature is good for matching or not. The quality of the feature is determined during training if a successful match is found within the first K neighbors in the dictionary. The classification branch is trained until its accuracy converges at $\approx 70\%$ with $K = 1$ neighbor. At test time the confidence score from the classification branch is used to compute a joint probability of all matched segments. During the last stage of the localization pipeline the candidates are sorted using the joint probability distribution and the pose is estimated using PROSAC [9] with applied geometric consistency constraints [7].

B. Learning the Segment Feature Space

To learn the feature space in the proposed architecture a variation of the triplet loss [26] is used. The triplet loss clusters together samples in feature space that have have been labelled as similar and tries to separate samples that have been labelled as different. By applying the same label to different examples of a segment we introduce invariance to many factors such as noise in the measurements or incomplete segments. The triplet loss defines a triplet as a combination of a sample (anchor) with other two samples, one with the same label and the other with a different label (Fig. 2, left). A pairwise term is defined as a combination of a sample (anchor) with a different sample that has the same label. To train our descriptor we use a variation of the triplet loss similar to the one defined in [27] - in our case we use a squared L_2 norm as shown in Eq. (1) while they use a regular L_2 norm. To train the model we use a batch size of 256 point clouds from which we extract a large set of triplet and pairwise terms.

$$\mathcal{L} = \mathcal{L}_{triplets} + \alpha \mathcal{L}_{pairs} + \lambda \|w\|_2^2. \quad (1)$$

During training we identify ≈ 300 times more triplets than pairwise terms. We balance this effect using the α parameter. In Eq. (1) w denotes the parameters of our model, weighted by $\lambda = 10^{-6}$. The proposed architecture contains only $300\,\mathrm{K}$ trainable parameters. We use an initial learning rate of $\eta = 10^{-3}$, decaying to minimum $\eta = 10^{-6}$ with ADAM [28] as optimizer. We describe the two losses in detail below.

The classification network is trained using the same data used for the embeddings layer. The features are precomputed and used as a dictionary. To train the classification branch, for each sample we find the K closest matches in the dictionary. A binary label is assigned for each sample: if a match is found in the dictionary it is labelled as 1, otherwise 0. The network optimizes a softmax cross-entropy loss for the classification branch, given the aforementioned labels.

C. Triplet Loss

As commented before, we use a variation of the definition of triplet loss similar to the one in [27]. This loss modified the original triplet loss to solve the vanishing gradient problem. We modify this loss by using squared L_2 norm, instead of the standard L_2 norm, as it created a better clustering of the feature space. The final triplet loss is defined as the sum of the following cost function over all the triplets:

$$\mathcal{L}_{triplets} = \sum_{(s_i, s_j, s_k) \in \mathcal{T}} c(s_i, s_j, s_k) \quad (2)$$

$$c(s_i, s_j, s_k) = \max\left(0, 1 - \frac{\|f_w(x_i) - f_w(x_k)\|_2^2}{\|f_w(x_i) - f_w(x_j)\|_2^2 + m}\right) \quad (3)$$

$\mathcal{T}$ denotes the set of all possible triplets. s_i and s_j are segments with the same label, while s_i and s_k are dissimilar. $f_w(x_i)$ is the output of the last descriptor layer for an input point cloud x_i and m is a margin regularizer. The latter denotes the minimum ratio for the Euclidean distances between dissimilar pairs of point clouds and similar ones. We use $m = 0.01$ in our experiments. In this manner the triplet loss will cluster similar objects together, while separating dissimilar ones to be farther apart in the feature space. The proposed architecture will provide a descriptor of the whole segment per point, centered at each of those points. Our feature is computed as the average of those descriptors. Our losses are computed using the averaged features. Fig. 2 (right) presents a PCA visualization of the feature space after the network has been trained with the corresponding segments. Different instances of the same tree are clustered together while other trees form their own clusters.

D. Pairwise Loss

The pairwise loss minimizes the distance between samples of the same class (Fig. 2, orange). As proposed by [27], we optimise the following:

ROC KITTI Sequence 00
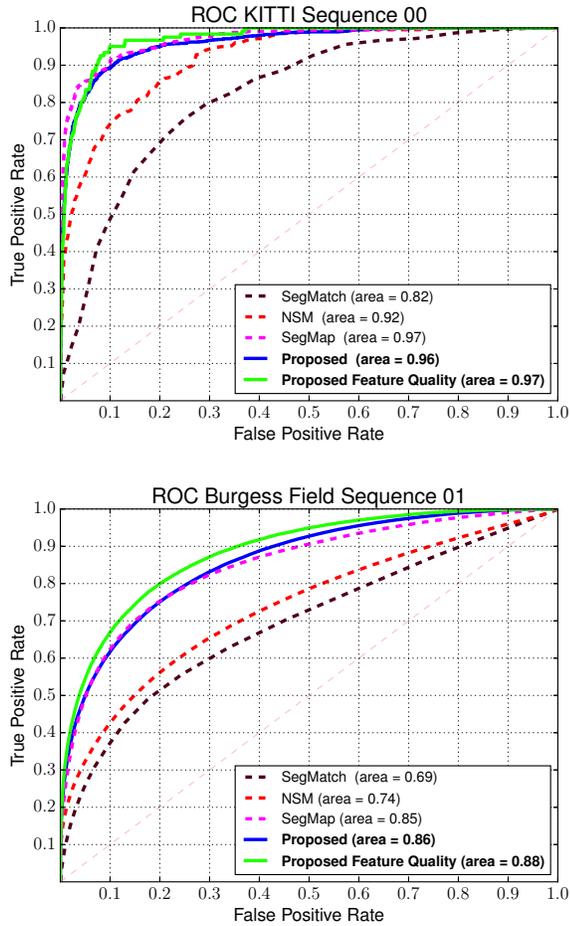

ROC Burgess Field Sequence 01

Fig. 3: Receiver Operating Characteristic (ROC) curves for the proposed and baseline approaches [5], [6], [7] on urban (top) and natural (bottom) datasets. The proposed approach performs comparably to the current state of the art on both environments. Learning approaches outperform hand-engineered ones in both urban and natural scenarios.

$$\mathcal{L}_{pairs} = \sum_{(s_i, s_j) \in \mathcal{P}} \| f_w(x_i) - f_w(x_j) \|_2^2 \qquad (4)$$

where $\mathcal{P}$ denotes the set of all pairs. The loss aids the training process by generating very tight clusters, which in turn improve the KNN retrieval during the matching phase.

IV. EXPERIMENTAL EVALUATION

The main focus of this work is to utilize learned features in point clouds to localize with respect to a prior map. The experiments are designed to support our key claims:

- A novel learned feature descriptor that generalizes across a variety of natural and urban datasets without the need for retraining.
- Real time operation on a CPU as a result of a compact network architecture.
- The approach produces a measure of quality for each feature which can be used during the pose estimation step to decrease computation

We provide comparisons against a popular method for segment-based localization, SegMatch, as proposed in [5], a data-driven incremental approach SegMap [6], and our previous work NSM [7]. As a metric to evaluate the approaches we have chosen to compare the True Positive Rate (TPR) against the False Positive Rate (FPR) for each classifier, the number of localizations on each dataset, and the computation time for each pipeline. A video accompanying the paper is available at http://ori.ox.ac.uk/esm-localization.

A. Datasets

We perform the evaluations on three datasets of our own which were collected in natural environments as well as on a publicly available dataset. First, we compared our novel descriptor against the baseline approaches using the KITTI [4] dataset (Karlsruhe, Germany), taken by a 3D Velodyne-HDL64 sensor in an urban scenario. We use Sequence 06 and 05 to train all descriptors, and Sequence 00 to test them. We also compare the localization performance on Sequence 00. Second, we compared localization performance in a park and forest datasets - George's Square (Edinburgh, Scotland) and Cornbury Park (Oxfordshire, UK). The former was captured by a Clearpath Husky UGV equipped with a SICK LMS511 LIDAR, while the latter was captured by a vehicle equipped with a Velodyne HDL32E. The datasets are described in detail in [7].

Finally, we used a dataset collected in a natural environment, located at Burgess Field (Oxfordshire, UK). The dataset was captured by a Clearpath Husky UGV equipped with a Bumblebee3 forward facing camera, a Velodyne VLP-16 LIDAR and a Push-broom LIDAR LMS151. The vehicle traversed the same loop of $\approx 2.8\,\text{km}$ twice a month during a span of six months. We used the Velodyne data of Sequence 02 and Sequence 01 for training and testing the classifiers. These two sequences were collected two weeks apart each other in February when the foliage was shed. We built a prior map from the Push-broom LIDAR of Sequence 02 using [2]. The source swathes were built using the VO method [2] for every 22 meters the vehicle traversed in Sequence 01.

To label the data we combined four consecutive Velodyne scans using the motion estimator from [2] and extracted segments using distance-based segmentation algorithm. All segments to be within $0.5\,\text{m}$ in consecutive point clouds are considered the same object and labelled as such. All segments further than $4.0\,\text{m}$ apart are considered a non-match. In this manner we extracted a total of 82871 segments across 38714 classes for training and 45409 segments in 20583 classes for testing. Training our model on Burgess Field takes 3 hours and 25 minutes from scratch. We have used this model to evaluate our approach on all natural datasets. For experiments carried out on urban datasets, we trained a different model on KITTI in 11 hours and 45 minutes. For all the experiments and methods we use Euclidean segmentation. We segmented the data into individual point clouds depending on the sampling density: 3D Velodyne data: $200 - 15000$ points due to higher frequency, Push-broom type LIDAR: $200 - 50000$ points due to higher density.
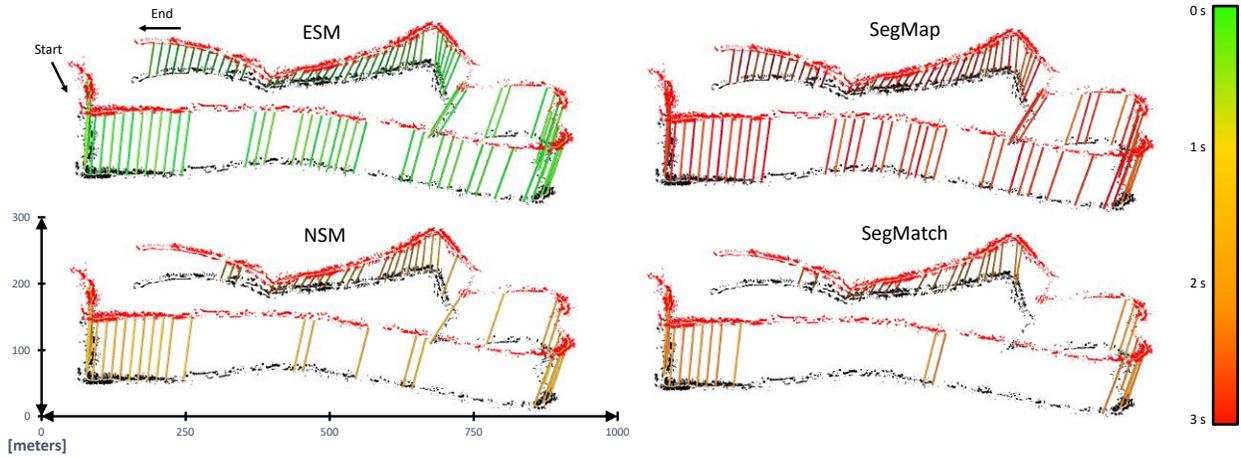
Fig. 4: Illustration of the localization performance of our approach (top left) compared to the baselines on the Burgess Field dataset. The current point clouds (red) are registered to a previous map (black). Each vertical line corresponds to a recognised place, while its color corresponds to the time taken to perform the localization.

Dataset Characteristics						Number of Localizations			
Name	Type	Length (meters)	Sensor	Num. Clouds		SegMatch	NSM	SegMap	ESM
KITTI	Urban	1300	Accum. Velodyne-HDL64E	134		41	53	62	**63**
George Square	City Park	500	Push-broom	30		8	9	8	**11**
Cornbury Park	Forest	700	Velodyne-HDL32E	1125		N/A	29	231	**248**
Burgess Field	Forest	2800	Push-broom	130		41	54	**95**	94

TABLE I: Number of localizations across one urban and three natural datasets. The proposed approach performs comparably to the state of the art. This experiments also shows the ability of ESM to handle different kinds of LIDAR sensors.

B. Classifier Performance

The first experiment evaluates the performance of our model, with and without the feature quality, on KITTI and Burgess Field datasets. Fig. 3 illustrates the Receiver Operator Characteristic (ROC) curves for each of the algorithms. A performance decrease is seen between the two datasets for all algorithms. We attribute this to the more challenging natural structure in the second environment. In brief, our learning approach generalizes better than engineered features (NSM, SegMatch) across datasets and performs comparably to the learned approach (SegMap). Pruning the matches based on the feature quality shows an improvement with respect to the basic features. For all models we have set the classifier thresholds at FPR=0.1 for urban and FPR=0.2 for forest environments. In addition, we did not retrain our approach for the each forested scenarios or the different sensor modalities. Showing that the proposed features can generalize between datasets and sensors.

C. Localization Performance

In the next set of experiments we aim to support our claim that the proposed algorithm performs comparably to the state of the art on both urban and natural environments when localizing, while requiring less computation. Tab. I shows the total number of loop closures detected in each of the datasets. For each algorithm we have optimized the parameters to retrieve the highest number of true localizations, while having zero false localizations. The learning approaches have not been retrained for each individual natural dataset but

still manage to detect about two times more loop closures than the engineered methods. In brief, our algorithm has similar accuracy to [6] on all datasets. Fig. 4 presents qualitative results for our approach compared to the baselines on the Burgess Field dataset, while also highlighting the computational times for each approach. In this environment the vehicle regularly detected loop closures during more than 2 km of operation. Fig. 5 demonstrates the performance of our algorithm across all the datasets. The method did not produce any false localizations, while operating in real time on a CPU.

D. Computation Efficiency

Finally, we evaluate the computational performance of the various approaches. We tested our approach on a mobile Intel Xeon E3-1535M CPU with the following configurations: (1) single-core of the processor, (2) multi-core on the same processor, (3) CPU + NVIDIA Titan Xp GPU. We have processed the Push-broom point cloud data from Burgess Field and recorded the mean computational speeds. The target map consisted of 2783 segments — individual trees, bushes, and interleaved shrubs. Individual source clouds were created for every 22 m travelled with each source cloud containing an average of 46 segments. In order to provide a fair comparison, we have utilized the same Euclidean segmentation method across algorithms, to create the same segments. This is important for the description stage of each pipeline. The segmentation was carefully designed to increase localization performance in this challenging natural

CPU Multi-core execution times								
Algorithm	**Desc. size**	**Segmentation**	**Preprocessing**	**Descriptor**	**Matching (K)**	**Prunning (RF)**	**Pose Est.**	**Total**
SegMatch	7 (647)	17	0	605	50 (200)	755	64	1491 ms
NSM	66	17	0	24	207 (200)	913	65	1226 ms
SegMap	64	17	15	5902	19 (25)	0	21	5974 ms
ESM	16	17	3	578	8 (25)	0	9	615 ms
CPU Single-core execution times								
Algorithm	**Desc. size**	**Segmentation**	**Preprocessing**	**Descriptor**	**Matching (K)**	**Prunning (RF)**	**Pose Est.**	**Total**
SegMatch	7 (647)	21	0	596	70 (200)	773	64	1524 ms
NSM	66	21	0	37	213 (200)	936	65	1272 ms
SegMap	64	21	26	25945	26 (25)	0	21	26039 ms
ESM	16	21	3	2126	12 (25)	0	11	2173 ms
GPU Computation comparison								
Algorithm	**Desc. size**	**Segmentation**	**Preprocessing**	**Descriptor**	**Matching**	**Pruning (RF)**	**Pose Est.**	**Total**
SegMap	64	17	15	15	73	19	15	154 ms
ESM	16	17	3	2	66	8	6	102 ms

TABLE II: Average computational times in milliseconds recorded per point cloud on the Burgess Field dataset.
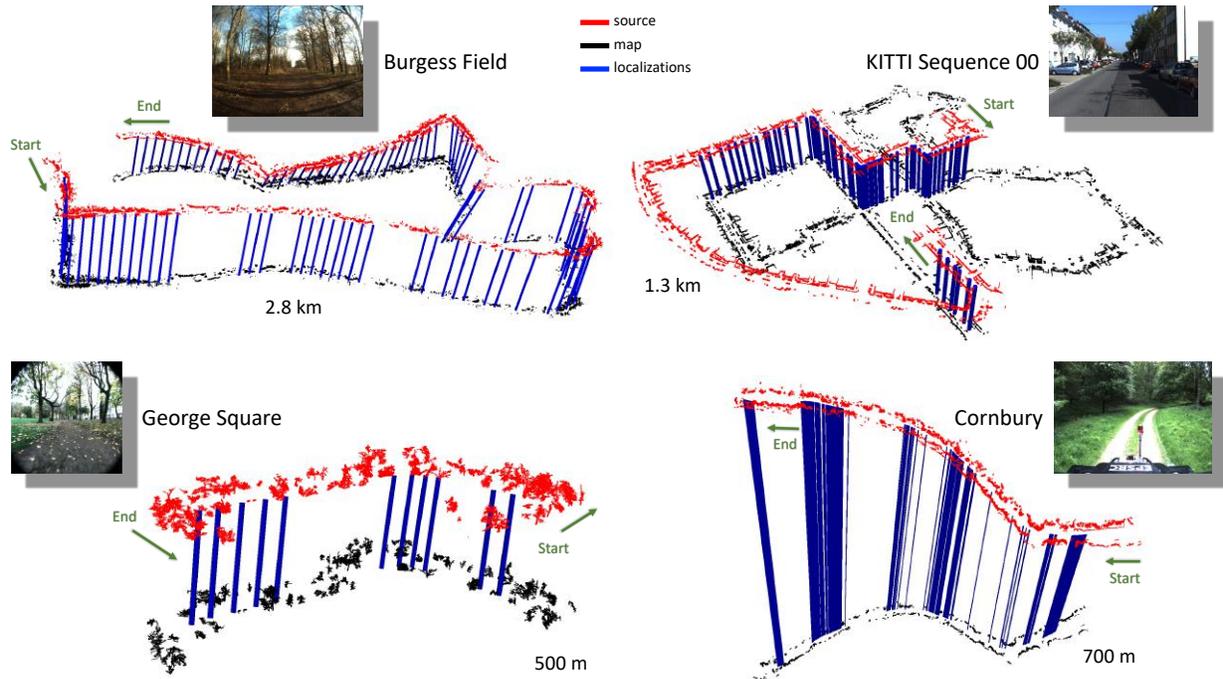

Fig. 5: Qualitative visualization of the performance of the proposed approach (ESM) on all datasets. The used datasets provide a variety of contexts to show the behaviour of ESM when coping with structured and unstructured locations.

environment. To estimate the pose we have used the same geometric consistency algorithm for all baselines substituting RANSAC [29] with PROSAC [9] for the proposed method to incorporate the quality of the features. In this setting, using an incremental geometric consistency as in [6] is not possible, as the Push-broom LIDAR only provides one observation of a segment.

We have empirically evaluated the number of neighbors in feature space (K) on Burgess Field to retrieve the most positive localizations, while keeping zero false positives. For this experiment K influences the KNN retrieval and Random Forest pruning speeds. $K = 25$ worked well for learning approaches as the features were very descriptive, while we kept $K = 200$ for NSM and SegMatch. Tab. II summarizes the average runtime performance, per point cloud, in mil-

liseconds. The size of the descriptor dimension corresponds directly to the computation time taken to describe a segment. The preprocessing and descriptor times scale linearly with the number of segments in a live point cloud (≈ 46 in Burgess Field). The time required for matching also depends on the embedding dimension. Point clouds tend to be of similar size, thus the computation time for the segmentation does not vary. The pose estimation depends on the number of segments and their closest neighbors (K). The total size of the SegMatch descriptor is 647 dimensions, only 7 of which are used during the matching stage. These 647 are compressed efficiently to the 45 processed by the Random Forest. This results in fast KNN retrieval and slower RF pruning. The total size of the NSM descriptor is 66, which are extended to 330 for the pruning stage, making it less effi-

cient for KNN retrieval and pruning. The SegMap descriptor requires PCA alignment and voxelization of the segments as preprocessing steps, after which the forward pass of the model is executed in C++. Even though the implementation is efficient, the $9.3\,\mathrm{M}$ parameters of the network describe a single segment in $0.5\,\mathrm{s}$ on a CPU. Tab. II also summarizes the recorded GPU times for the learned approaches with $K = 25$ during matching.

We focused our analysis on CPU-only solutions due to the particular efficency of our approach. Our network has only $300\,\mathrm{K}$ parameters which allows the descriptor to work in real time on a CPU. Also, our embedding dimensions are small making the matching quick.

V. Conclusion

In this paper, we presented a novel approach to place recognition based on LIDAR segment matching in both urban and natural environments. Our method exploits an efficient deep learning architecture that operates directly on point cloud data without the need of extensive preprocessing. This allows us to successfully detect loop closures at $\approx 1\,\mathrm{Hz}$ while using only a CPU. We implemented and evaluated our approach on four different datasets containing natural forest and parkland as well as urban scenes and provided comparisons to other existing techniques. Our approach operates in real time on a CPU while achieving matching performance comparable to the state of the art, which requires a GPU to run in real time. The experiments suggest that our approach can be applied to mobile robots with limited computational power. In future work we are interested in deploying the approach on an ANYMAL quadruped and a UAVs both of which lack a GPU.

VI. Acknowledgements

We would like to thank Oliver Barlett and the rest of the team at the Oxford Robotics Institute for the datasets.

References

[1] P T Furgale and T D Barfoot. Visual Teach and Repeat for Long-Range Rover Autonomy. *Journal of Field Robotics,* special issue on "Visual mapping and navigation outdoors", 27(5):534–560, 2010.
[2] Winston Churchill and Paul Newman. Experience-based Navigation for Long-term Localisation. *Intl. J. of Robotics Research*, 2013.
[3] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M. Seitz, and Richard Szeliski. Building Rome in a Day. *Communications of the ACM*, 54(10):105–112, 2011.
[4] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *Intl. J. of Robotics Research*, 32(11):1231–1237, 2013.
[5] Renaud Dubé, Daniel Dugas, Elena Stumm, Juan Nieto, Roland Siegwart, and Cesar C. Lerma. SegMatch: Segment based place recognition in 3D point clouds. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 5266–5272, 2017.
[6] Renaud Dubé, Andrei Cramariuc, Daniel Dugas, Juan Nieto, Roland Siegwart, and Cesar Cadena. SegMap: 3D Segment Mapping using Data-Driven Descriptors. In *Robotics: Science and Systems (RSS)*, 2018.
[7] Georgi Tinchev, Simona Nobili, and Maurice Fallon. Seeing the Wood for the Trees: Reliable Localization in Urban and Natural Environments. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2018.
[8] Yangyan Li, Rui Bu, Mingchao Sun, and Baoquan Chen. PointCNN. *arXiv preprint arXiv:1801.07791*, 2018.

[9] O. Chum and J. Matas. Matching with PROSAC - Progressive Sample Consensus. In *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, 2005.
[10] Fan Haoqiang, Hao Su, and Leonidas J Guibas. A Point Set Generation Network for 3D Object Reconstruction from a Single Image. In *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, 2017.
[11] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Conf. on Neural Information Processing Systems (NIPS)*, 2017.
[12] Sebastian Thrun, Dieter Fox, Wolfram Burgard, and Frank Dellaert. Monte Carlo Localization for Mobile Robots. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 1999.
[13] Jesse Levinson, Michael Montemerlo, and Sebastian Thrun. Map-Based Precision Vehicle Localization in Urban Environments. In *Robotics: Science and Systems (RSS)*, 2007.
[14] M. W. M. Gamini Dissanayake, Paul Newman, Steven Clark, Hugh F. Durrant-whyte, and M. Csorba. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Trans. Robotics*, 17:229–241, 2001.
[15] Ryan W Wolcott and Ryan M Eustice. Robust LIDAR localization using multiresolution Gaussian mixture maps for autonomous driving. *Intl. J. of Robotics Research*, 36:292–319, 2017.
[16] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-Time Loop Closure in 2D LIDAR SLAM. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2016.
[17] François Pomerleau, Francis Colas, Roland Siegwart, and Stéphane Magnenat. Comparing ICP variants on real-world data sets. *Autonomous Robots*, 34(3):133–148, April 2013.
[18] B. Douillard, A. Quadros, P. Morton, J. P. Underwood, M. De Deuge, S. Hugosson, M. Hallstrm, and T. Bailey. Scan segments matching for pairwise 3D alignment. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2012.
[19] Christopher Zach, Adrian Penate-Sanchez, and Minh-Tri Pham. A dynamic programming approach for fast and robust object pose recognition from range images. In *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, 2015.
[20] Gil Elbaz, Tamar Avraham, and Anath Fischer. 3D Point Cloud Registration for Localization using a Deep Neural Network Auto-Encoder. In *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, pages 2472–2481, 2017.
[21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Conf. on Neural Information Processing Systems (NIPS)*, 2012.
[22] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3D-R2N2: A Unified Approach for Single and Multi-view 3D Object Reconstruction. In *Eur. Conf. on Computer Vision (ECCV)*, 2016.
[23] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, 2017.
[24] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
[25] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *Intl. Conf. on Learning Representations (ICLR)*, 2016.
[26] Kilian Q Weinberger, John Blitzer, and Lawrence K. Saul. Distance Metric Learning for Large Margin Nearest Neighbor Classification. In *Conf. on Neural Information Processing Systems (NIPS)*, 2006.
[27] Paul Wohlhart and Vincent Lepetit. Learning descriptors for object recognition and 3D pose estimation. In *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, pages 3109–3118, 2015.
[28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Intl. Conf. on Learning Representations (ICLR)*, 2014.
[29] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.