



This is a repository copy of *NDT-Transformer : large-scale 3D point cloud localisation using the normal distribution transform representation*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/172335/>

Version: Accepted Version

Proceedings Paper:

Zhou, Z., Zhao, C., Adolfsson, D. et al. (4 more authors) (2021) NDT-Transformer : large-scale 3D point cloud localisation using the normal distribution transform representation. In: Proceedings of the 2021 IEEE International Conference on Robotics and Automation. ICRA 2021 - IEEE International Conference on Robotics and Automation, 30 May - 05 Jun 2021, Xi'an, China. Institute of Electrical and Electronics Engineers , pp. 5654-5660. ISBN 978-1-7281-9077-8

<https://doi.org/10.1109/ICRA48506.2021.9560932>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

NDT-Transformer: Large-Scale 3D Point Cloud Localisation using the Normal Distribution Transform Representation

Zhicheng Zhou^{1†}, Cheng Zhao^{2†}, Daniel Adolphsson³, Songzhi Su⁴,
Yang Gao⁵, and Tom Duckett⁶, Li Sun^{1*}

Abstract—3D point cloud-based place recognition is highly demanded by autonomous driving in GPS-challenged environments and serves as an essential component (i.e. loop-closure detection) in lidar-based SLAM systems. This paper proposes a novel approach, named NDT-Transformer, for real-time and large-scale place recognition using 3D point clouds. Specifically, a 3D Normal Distribution Transform (NDT) representation is employed to condense the raw, dense 3D point cloud as probabilistic distributions (NDT cells) to provide the geometrical shape description. Then a novel NDT-Transformer network learns a global descriptor from a set of 3D NDT cell representations. Benefiting from the NDT representation and NDT-Transformer network, the learned global descriptors are enriched with both geometrical and contextual information. Finally, descriptor retrieval is achieved using a query-database for place recognition. Compared to the state-of-the-art methods, the proposed approach achieves an improvement of 7.52% on average top 1 recall and 2.73% on average top 1% recall on the Oxford Robotcar benchmark.

I. INTRODUCTION

There is a high demand for point-cloud-based navigation systems that are able to robustly localise robots or autonomously driving cars in GPS-challenging environments. Point-cloud-based global localisation can also be employed as a loop-closure detection module, which is an essential component of simultaneous localisation and mapping (SLAM) systems. A practical method is to use GPS to acquire the coarse global location and point cloud registration methods, such as ICP, to obtain a more accurate pose estimate. However, GPS is not always available and reliable, so that alternative solutions using only sensory data are required.

The conventional point cloud registration methods and particle filtering methods do not scale well to large-scale environments. The main challenges for current robot localisation systems are generality and scalability. For example, model-based methods [1], [2], [3] can learn the 6DOF global pose as a regression model and infer the pose in real-time. Sadly, however, regression-based methods do not generalise to novel environments. SegMatch-like methods [4], [5], [6],

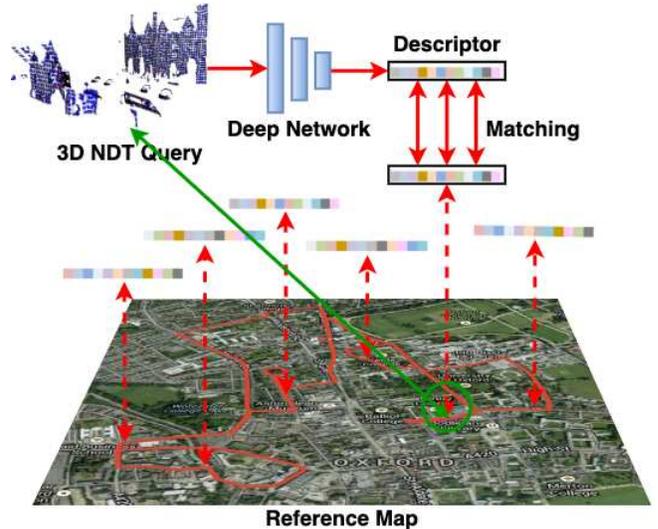


Fig. 1: 3D point (NDT)-based loop-closure detection in a large-scale environment (Oxford city centre).

[7] register two point cloud maps on intermediate representations (e.g. semantic segments), which has proven to be robust but not computationally efficient.

This paper addresses large-scale localisation as a large-scale point cloud retrieval problem, where the continuous robot car trajectories are discretized into consecutive ‘places’. Inspired by the success of the Transformer model [8] in natural language processing, we propose a novel approach – *NDT-Transformer* – to transform the point cloud into a ‘long sentence’ of geometry-enriched NDT cells and then to a site-specific feature signature. During inference, only the descriptor of the query data is required to be computed online, while the other descriptors in the database can be computed once offline and stored in memory, thus enabling real-time topological localisation in large-scale environments. Consequently, both scalability and generalisation are achieved.

The main contributions can be summarised as follows: 1) A computationally efficient method is proposed for large-scale point-cloud-based localisation where NDT is used as an intermediate representation. Our approach condenses a dense point cloud into a lightweight representation with maximal preservation of the geometrical features; 2) A novel network architecture, named *NDT-Transformer* is devised to learn a global descriptor with contextual clues from a set of 3D NDT cell representations; 3) The proposed method achieves

†Authors contribute equally.

This project is funded by EPSRC FAIR-SPACE Hub (EP/R026092/1) and EU Horizon 2020 ILIAD (No. 732737).

*The corresponding author. li.sun@sheffield.ac.uk

¹Visual Computing Group, University of Sheffield, UK

²Department of Engineering Science, University of Oxford, UK

³Mobile Robotics and Olfaction Lab, Orebro University, Sweden

⁴IMT Lab, Xiamen University, China

⁵STAR LAB, Surrey Space Centre, University of Surrey, UK

⁶Lincoln Centre for Autonomous Systems, University of Lincoln, UK

the state-of-the-art performance in point-cloud-based place recognition, which can be an important supplement for NDT-based SLAM and Monte-Carlo localisation methods.

The processed data, code and trained model can be found on our project website: <https://github.com/dachengxiaocheng/NDT-Transformer.git>

II. RELATED WORK

A. Point cloud-based global localisation

The state-of-the-art point cloud based localisation methods can be categorised into three main streams: regression-based [1], [2], [3], intermediate-representation-based [4], [5], [6] and global-feature-based [9], [10], [11], [12], [13], [14], [15].

Following the success of deep pose estimation in image-based global localisation, some methods [1], [2], [3] propose to use deep regression networks to learn the 6DOF global pose. These model-based approaches are extremely efficient but not generalisable. In other words, a new model is required to be trained for new environments.

Intermediate-representation-based methods [4], [5], [6], [7] first segment the map into intermediate representations (i.e. semantic segments) then register the segmented parts, rather than directly registering the point clouds. The intermediate parts can be described as distinctive features [4] or deep-learned features [5], [6], [7] may be used to eliminate the false positives. Though the matching efficiency can be improved by using intermediate representations, significant running time is spent on cloud segmentation and real-time performance cannot be guaranteed. Therefore, these methods are mostly used for loop-closure detection rather than large-scale localisation.

The mainstream methods for point-cloud-based localisation are based on global features, where a generalisable feature extractor is designed or learned to obtain the place signature for retrieval. Before deep learning began dominating the machine learning community, hand-designed heuristics features such as M2DP [9], Scan Context [10] and DELIGHT [11] were well-studied to represent the 3D point cloud for the localisation. In terms of learning-based methods, the Scan Context Image-based network [12] and OverlapNet [13] convert the 3D lidar scan to a 2D image according to geometric knowledge, then deploy a 2D convolution-like network to learn a representation.

Instead of using hand-crafted features, PointNetVLAD [16] combines PointNet [17] and NetVLAD [14] to learn a global descriptor based on metric learning. However, the PointNet-like architecture ignores the spatial distribution and contextual cues within the 3D point cloud. Extracting efficient contextual information from the irregular 3D point cloud is another important challenge for 3D loop-closure detection. The following work LPD-Net [15] employs a classical DGCNN [18]-like network to enhance the feature descriptor by KNN-based aggregation in both feature space and Cartesian space. They also introduce 10 different kinds of geometric local features to feed the network during training, achieving the state-of-the-art (SOTA) performance. Most of these point-cloud-based

methods require a downsampling operation on the original 3D dense 3D point cloud due to the GPU size limitation. However, this downsampling operation can lead to a loss of local geometric information. Lastly, due to the real-time attribute of SLAM, one challenging aspect of 3D loop-closure detection is the real-time performance requirement, which means that lightweight networks are required.

B. NDT-based localisation

The Normal Distribution Transform (NDT) is a classic method to represent a 2D or 3D point cloud as differential multi-variant Gaussian distributions for 2D laser [19] and 3D lidar mapping [20], [21]. Compared to point cloud or grid-map based representations, NDT-based relocalisation, i.e. NDT-MCL [22], demonstrates advanced localisation precision and repeatability, as the NDT is an inherently probabilistic and geometrical representation to model the likelihood. In other words, the structural and geometrical information of the map is implicitly interpreted by the NDT cell parameters. Benefiting from this characteristic, NDT-based localisation can condense the bulky point cloud map to memory efficient NDT cells to make the localisation scalable. Submap-based NDT [23] can eliminate the uncertainties caused by different robot perspectives, and further improve the localisation accuracy and scalability in large-scale and long-term applications [2]. Moreover, NDT can be used to provide place signatures for loop closure, e.g. using NDT histograms [24] and semantic-NDT histograms [25], semi-supervised place categorisation [26], and to create interest point descriptors [27] for efficient registration or localisation.

III. METHODOLOGY

A. Problem Formulation

Given a 6DoF trajectory of the robotcar and the synchronised lidar scans, submaps $M = \{m_1, \dots, m_N\}$ can be built by dividing the trajectory into fixed range intervals and mapping the point cloud into a local reference frame. Then the obtained submap can be further represented as NDT cells. Specifically, the space of a submap will be uniformly divided into grid cells, and the 3D points within a cell will be used to estimate the NDT parameters. Practically, the number of cells varies greatly due to the different density of each submap. A spatially distributed sampling filter \mathcal{G} is applied to guarantee that the numbers of cells in all submaps are the same, i.e. $|\mathcal{G}(m_1)| = \dots = |\mathcal{G}(m_N)|$. A deep model called NDT-Transformer is proposed to learn a function $f(\cdot)$ which represents the input NDT cell representations as a fixed-size global descriptor $f(\mathcal{F})$ where $\mathcal{F} = \mathcal{G}(m)$. Then the Euclidean distance function $d(\cdot)$ can be adopted to measure the similarity, i.e. $d(f(\mathcal{F}), f(\mathcal{F}_{pos})) < d(f(\mathcal{F}), f(\mathcal{F}_{neg}))$, where \mathcal{F} is similar to \mathcal{F}_{pos} and dissimilar to \mathcal{F}_{neg} .

Hence the problem of 3D loop closure detection is illustrated as follows. Denoting the query submap as m_q , the task aims at searching the database M for the most similar submap m_* . The problem can be resolved by searching for the nearest neighbour of a submap $m_* \in M$ in the feature space, i.e. $f(\mathcal{F}_{m_*})$, which has the smallest Euclidean

distance to the query point cloud (feature) $f(\mathcal{F}_{m_q})$. That is, $d(f(\mathcal{F}_{m_q}), f(\mathcal{F}_{m_*})) < d(f(\mathcal{F}_{m_q}), f(\mathcal{F}_{m_i}))$, $\forall m_i \in M, i \neq *$.

An example of 3D place recognition is illustrated in Fig. 1. The red line on the map shows the entire route navigated by an autonomous vehicle in the city centre of Oxford, UK. A set of point clouds for this run can be constructed by combining the scans with relative global poses. The objective is to identify the nearest match of the query point cloud from another run. This paper represents all point clouds as a fixed number of 3D NDT cells (as shown in Fig 2) and utilises the NDT-Transformer to convert them to site-specific global descriptors. Moreover, only the query descriptor is required to be computed online, while the other descriptors in the database are computed once offline and stored in memory, which can guarantee the real-time performance.

B. 3D NDT Representation

In contrast to most existing work, we represent the point cloud submap using the 3D Normal Distributions Transform (NDT) [21] rather than downsampling the 3D point cloud to a fixed number of points. As illustrated in Fig. 2, the 3D-NDT representation is a compact ‘spherical’ structure, which can not only preserve the geometrical shape of a group of local 3D points but can also significantly decrease the memory complexity. In other words, we can use very memory efficient resolutions such as $(1\text{m} \times 1\text{m} \times 1\text{m})$ without losing too much geometrical information. The normal distribution $\mathcal{N}(\mu, \mathbf{C})$ of each NDT cell consists of the mean vector μ and the covariance matrix \mathbf{C} defined as,

$$(\mu, \mathbf{C}) = \left(\frac{1}{n} \sum_{k=1}^n \mathbf{x}_k, \frac{1}{n-1} \sum_{k=1}^n (\mathbf{x}_k - \mu)(\mathbf{x}_k - \mu)^T \right). \quad (1)$$

where $\mathbf{x}_{k=1, \dots, n}$ are 3D points in each cell.

We employ a spatially distributed sampling approach \mathcal{G} to generate evenly distributed NDT cells,

$$\mathcal{F} : m^s \rightarrow \mathcal{G}([\mu_i, \mathbf{C}_i]_{i=1..k}), \quad (2)$$

where m^s refers to submap of size s and \mathcal{G} refers to NDT representation of size k cells.

The spatially distributed sampling can reduce the point density from an arbitrary input size s to a fixed number of points k and simultaneously preserve structural information from the original submap. This operation consists of three steps. Firstly, we downsample the original submap: m^s of size s to $m^{k*c \pm t\%}$ (within some tolerance factor t), where $c > 1.0 + t$ is an oversampling factor that makes sure that an excessive amount of points are produced. The downsampling is done by performing binary search to find the voxel size that produces $k * c \pm t$ points. Secondly, we estimate NDT distributions \mathcal{F} with the downsampled points as centroids. That is, for each point in the downsampled submap $\mathbf{q}_i \in m^{k*c \pm t\%}$, an NDT cell (μ_i, \mathbf{C}_i) is computed from all points within a radius r of \mathbf{q}_i in the original submap m^s . The final step removes cells with the highest mutual information until k cells remain. This ensures that the structural information is

maintained. Mutual information is computed from the symmetric *Kullback–Leibler* divergence $1/2 * (D_{KL}(P||Q) + D_{KL}(Q||P))$. The neighbour with the lowest *KL*-divergence has the highest mutual information.

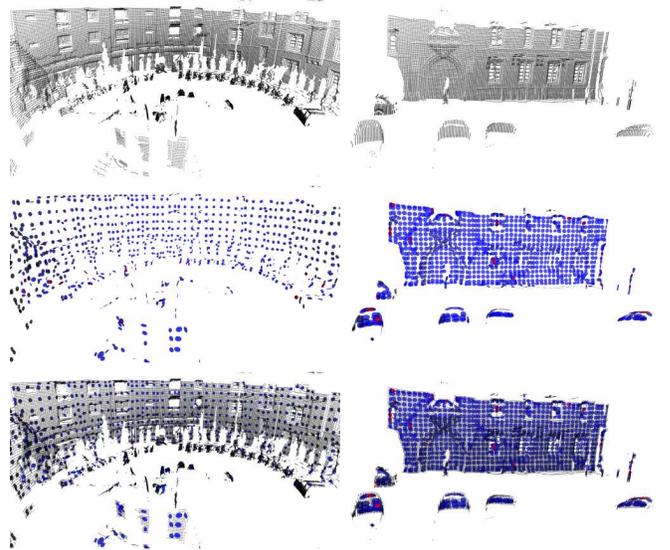


Fig. 2: First row: dense 3D point cloud, second row: 3D NDT cells, third row: 3D NDT cells within the dense 3D point cloud.

C. NDT-Transformer Network

After converting the 3D submap m to 3D NDT representation \mathcal{F} , this representation is fed into a NDT-Transformer network f to obtain a descriptor ξ of the 3D submap. That is, $\xi = f(\mathcal{F})$ where $\mathcal{F} = \mathcal{G}(m)$. As shown in Fig. 3, the NDT-Transformer network consists mainly of three components: 1) NDT representation module that consists of a point transform and uncertainty backpropagation, 2) A residual transformer encoder, and 3) Net Vector of Locally Aggregated Descriptors (NetVLAD).

In order to achieve rotational invariance of the descriptor, the point transform and uncertainty backpropagation are applied to each original NDT representation $\mathcal{F} = (\mu, \mathbf{C})$. Firstly, a transition matrix $T_{3 \times 3}$ is learned through a Transform Net (T-Net) [17] from the NDT mean vectors μ . Then a new NDT representation $\mathcal{F}_T = (\mu_T, \mathbf{C}_T)$ can be obtained by assembling the transformed points and propagated uncertainties,

$$\mathcal{N}(\mu_T, \mathbf{C}_T) = \mathcal{N}(T_{3 \times 3} \mu, T_{3 \times 3} \mathbf{C} T_{3 \times 3}^T), \quad (3)$$

where $T_{3 \times 3} = \mathcal{T}(\mu)$ and \mathcal{T} refer to the T-Net. Feature aggregation on irregular data such as NDT cells is still a challenging problem. PointNetVLAD [16] utilises PointNet [17] to transform the points’ position information to a high-dimensional feature representation. However, this representation of each point suffers from the ambiguities due to a lack of contextual cues. Although local feature aggregation can be achieved via mean and covariance calculation within each NDT cell, there is no aggregation

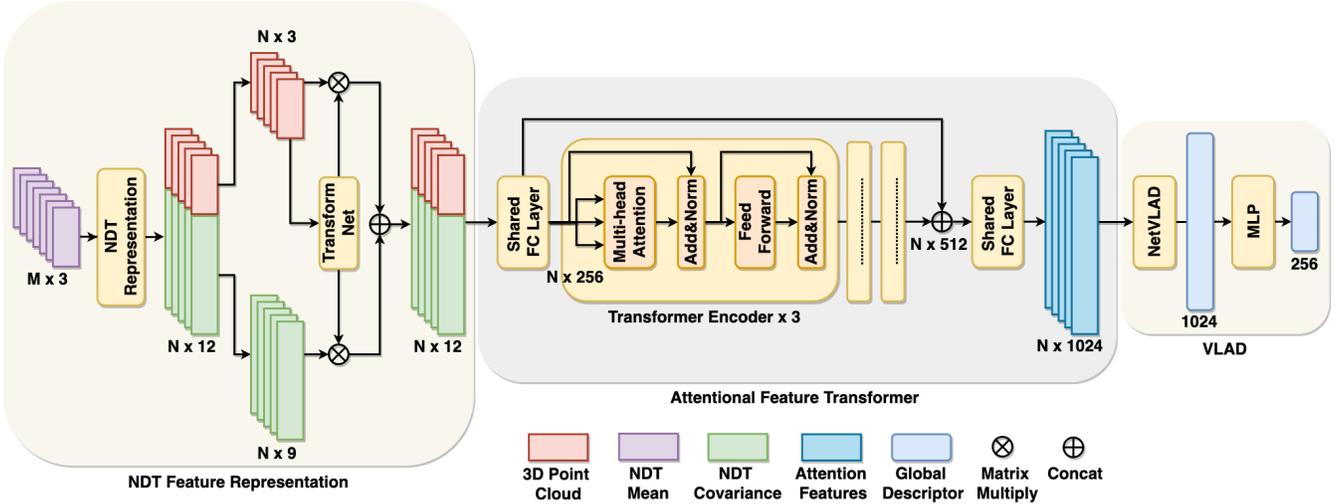


Fig. 3: NDT-Transformer Network Architecture

between the individual NDT representations. Our intuition is to employ an attention mechanism to learn the underlying context between landmarks (NDT cells). In contrast to most existing approaches [15], [18], which utilise kNN grouping for contextual information modelling, we employ a residual transformer encoder to aggregate the contextual cues for one NDT cell from the other NDT cells to increase the distinctiveness of the representation.

As shown in the middle part of Fig. 3, the network includes three transformer encoder stacks in series, and two shared linear stacks at the head and bottom with a shortcut skip connection. Each Transformer encoder is composed of a series of modules, i.e. multi-head self-attention (MHSA), feed-forward network (FFN) and layer normalization (LN), which can be stacked on top of each other multiple times. The transformer encoder can learn co-contextual information/message *Attn* captured by a self-attention mechanism,

$$Attn([Q_i, K_i, V_i]) = \text{concat}([\text{softmax}(\frac{Q_i \cdot K_i^T}{\sqrt{d_k}})V_i]), \quad (4)$$

where d_k denotes the dimension of queries and Q_i, K_i, V_i stand for i th head of queries, keys, values of the NDT cell representation, respectively. In our implementation, four-head attention (i.e. $i = 1, 2, 3, 4$) is used to enhance the discriminativeness of the feature attributes. The self-attention mechanism can automatically build connections between the current NDT cell and the other salient interesting NDT cells. The attentional representation Φ can be obtained as,

$$\Phi' = LN(\mathcal{F}_T + Attn), \quad (5)$$

$$\Phi = LN(FFN(\Phi') + \Phi'). \quad (6)$$

By this means, the local representation \mathcal{F}_T of each NDT cell is upgraded to an attentional representation Φ .

Following PointNetVLAD [16], we choose NetVLAD [14] instead of a max-pooling layer to improve permutation invariance for the descriptor of the 3D point cloud. NetVLAD is designed to aggregate a set of local descriptors and

generate one global descriptor vector. It records statistical information with respect to local signatures and sums the differences between these signatures and their respective cluster. In contrast to conventional VLAD, the parameters of NetVLAD, especially the assignment score, are learnt during training in an end-to-end manner. After going through the NetVLAD followed by a Multi-Layer Perceptron (MLP), a set of NDT cell representations/descriptors Φ can be fused to obtain a fixed-size global descriptor vector ξ to describe the 3D point cloud,

$$\xi = MLP(NetVLAD(\mathcal{F}_T \oplus \Phi)), \quad (7)$$

where \oplus refers to the concatenation.

D. Metric Learning

Following the PointNetVLAD architecture [16], the Lazy Quadruplet loss is used to achieve metric learning between the query global descriptor ξ_q and the positive ξ_{pos} , negative ξ_{neg} , and hard negative ξ_{neg*} examples, which are randomly picked during training from different locations to ξ_q . The Lazy Quadruplet loss is:

$$\mathcal{L}_Q(\xi_q, \xi_{pos}, \xi_{neg}, \xi_{neg*}) = \max([\alpha + d(\xi_q, \xi_{pos}) - d(\xi_q, \xi_{neg})]_+, \max([\beta + d(\xi_q, \xi_{pos}) - d(\xi_q, \xi_{neg*})]_+)) \quad (8)$$

where α, β are constant margin values, and $d(\cdot, \cdot)$ is the Euclidean distance. The core idea of Quadruplet loss is to minimise the distance between the query and positive global descriptors, and maximise the distance between the query and negative global descriptors, while simultaneously keeping a proper distance between all negative descriptors (i.e. ξ_{neg} and ξ_{neg*}).

IV. EXPERIMENT

A. Dataset Benchmark

The Oxford Robotcar dataset [28], which provides long-term multi-session data, was used to evaluate the proposed method. In this dataset, a Sick LMS-151 2D lidar scanner

mounted on the rear of the vehicle is used for mapping. Each scan contains a set of triplet records (x, y, R) , where the former two refer to coordinates in 2D Cartesian space and the latter refers to lidar infrared reflectance information. The dataset provides an accurate 6DoF trajectory of the travelled urban environment through fusing GPS with inertial sensors, and a dense point cloud map can be built by transforming 2D lidar scans with respect to the global poses. Then the places are generated with a fixed interval (10m for training and 20m for evaluation) and the global 3D map is divided into a set of submaps using a fixed length trajectory segment (20m).

In this experiment, we use 44 runs for training and 22 runs for testing from the Oxford RobotCar dataset. We obtain a total of 21898 submaps for training and 3071 submaps for testing. For the NDT cell generation, we use 2000 cells with a resolution of 0.8m as the input to We follow the definition of query, positive and negative data-pairs in PointNetVLAD [16], that is, the places within 10m to the query are classified as positives while places beyond 50m to the query are classified as negatives via KD-Tree search.

B. Neural Network Training

Similar to PointNet [17], an input Transform Network (T-Net) is required to achieve rotational invariance by transforming the input map to a canonical view. The hyperparameters of T-Net are the same as PointNet [17]. The two shared linear stacks before and after transformer encoders consists of linear layers with hidden variables with 256 and 1024 separately followed by Batch Normalisation layer and ReLU activation layer. The three Transformer encoders have the same hyperparameter settings: the dimensions of the input and output are set to 256, the dimension of the feed-forward layer is set to 1024, the dropout rate is 0.1, and the number of heads is set to 4. For the hyperparameter settings within NetVLAD, the dimensions of the input and output are set to 1024 and 256, respectively, and the size of the cluster is set to 64.

For the metric learning, the network is trained with 20 epochs with a batch size of 2. There are 1 query submap, 2 positive submaps, 18 negative submaps, and 1 hard negative submap in one single batch. Margin values of $\alpha = 0.5$ and $\beta = 0.2$ are used. The multi-step learning policy is used and the learning rate decay is a fixed value of 0.1 applied on epoch 9 and 15. The initial learning rate is $1e-5$ and the momentum is a fixed value of 0.9. The network is implemented in Pytorch and trained on a tower with Intel i9 CPU and two NVIDIA RTX Titan GPUs accelerated by CUDA and cuDNN.

C. Performance Evaluation

Following PointNetVLAD [16] on the Oxford Robotcar dataset benchmark [28], the recall indices, i.e., the Average Recall@N (N=1,2,3...25) and Average Recall@1% are employed as metrics for the performance evaluation of 3D loop closure detection.

NDT-Transformer achieves 93.80% top 1 (@1) average recall and 97.65% top 1% (@1%) average recall, respectively. As shown in Table I, ND-Transformer achieves a significant improvement compared to PointNetVLAD [16] with an improvement of 31.04% on top 1(@1) average recall and an improvement of 16.64% on top 1% (@1%) average recall. The performance of NDT-Transformer is superior to the SOTA, LPD-Net [15], with 7.52% improvement for top 1(@1) average recall and 2.73% improvement for top 1% (@1%) average recall. Although LPD-Net utilises 10 different kinds of local features, the NDT-Transformer only uses the NDT representations. The average recall curves from top 1 (@1) to top 25 (@25) of PointNetVLAD, LPD-Net and NDT-Transformer are provided in Fig 5.

The ablation studies for NDT-Transformer are provided in Table II. The average recall of top 1 (@1) and top 1% (@1%) decrease 17.87% and 9.76%, respectively, if only the point position without covariance matrices is used for learning. The average recall of top 1 (@1) and top 1% (@1%) show a decrease of 6.92% and 3.2%, respectively, if only the covariance matrices (i.e. without point position) are used for the network training. If the T-Net is excluded from the network, the average recall of top 1 (@1) and top 1% (@1%) decrease by 1.98% and 1.05% due to the absence of the point transform and uncertainty propagation for the NDT representation. Average recalls from top 1 (@1) to top 25 (@25) of the studied ablation baselines are provided in Fig. 5.

Network	Ave recall @1	Ave recall @1%
PointNetVLAD[16]	62.76	81.01
LPD-Net [15]	86.28	94.92
NDT-Transformer	93.80	97.65

TABLE I: The comparison result of PointNetVLAD, LPD-Net and NDT-Transformer, where average recall (%) at top 1 (@1) and the top 1% (@1%) are reported.

Network	Ave recall @1	Ave recall @1%
NDT-Transformer-P	75.93	87.89
NDT-Transformer-C	86.88	94.45
NDT-Transformer-noT	91.82	96.60
NDT-Transformer	93.80	97.65

TABLE II: The ablation study of NDT-Transformer is given. -P refers to our method using 3D points only. -C refers to our method using covariance matrices only. noT is the network architecture without using T-Net.

Finally, we show some examples of retrieval results on the Oxford Robotcar dataset using NDT-Transformer, where the ranking of correctly retrieved candidates is plotted in different colours. As shown in Fig. 4, the first row images are database submaps randomly picked from 22 tests in the experiment, in which the red trajectories are queries from previously unseen areas. The second row shows the heat maps indicating the retrieval results. Colours from red to yellow represent correctly recognised places in the top 1 to

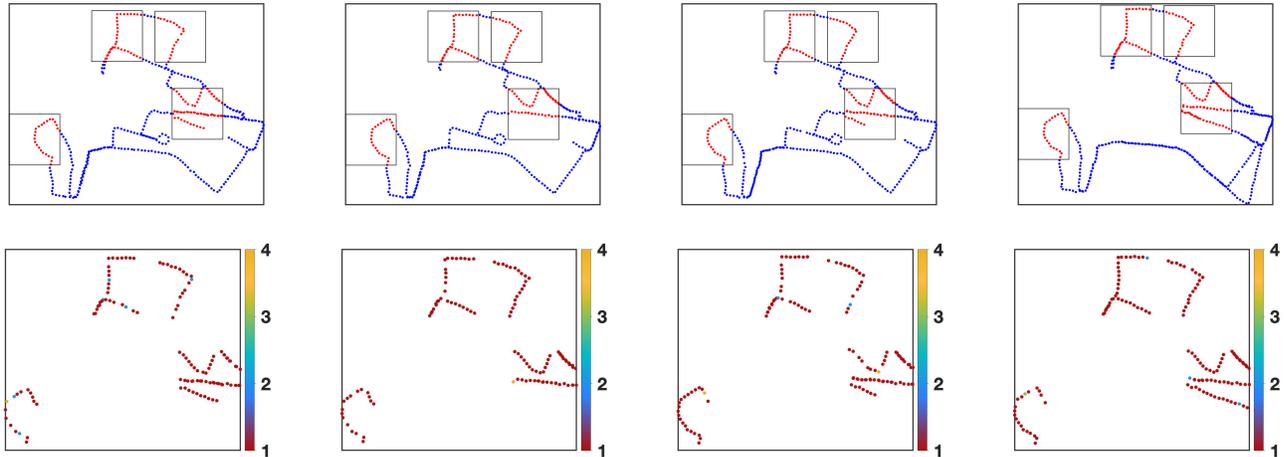


Fig. 4: In this figure, some examples of the database-query results of the proposed NDT-Transformer are shown. Each column refers to an experiment randomly picked in our evaluation (i.e. date 2014-11-14, 2014-12-09, 2015-02-03, 2015-04-24, on Oxford Robocar dataset). The first row shows the database locations where red dots are testing areas (i.e. streets never seen before). The second row shows the retrieval results (the ranking of successfully retrieved candidates)

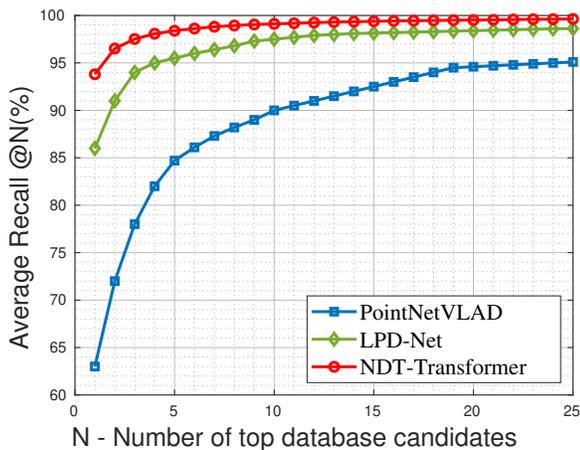


Fig. 5: The average recall curves from top 1 (@1) to top 25 (@25) of PointNetVLAD, LPD-Net and NDT-Transformer.

4 candidates. The KD-Tree is employed to search the top N matched candidates from the database. The figure shows that most of the queries are paired with the correct database locations as the first ranking candidate. It worth noting that the query areas (shown as four rectangles) are previously unseen streets, which means that no data from these areas were used in the training set. Therefore, this experimental result shows that our deep-learned features can be used for localisation in unseen areas (novel locations). Moreover, the inference time of our approach averages 0.034 seconds per point cloud (0.03 seconds for NDT conversion and 0.004 seconds for the transformer network) and the computational complexity of retrieval is $\mathcal{O}(\log n)$.

V. CONCLUSION

This paper proposes a novel approach, NDT-Transformer, to interpret a dense point cloud submap and provide a distinctive and discriminative location descriptor, and addresses the

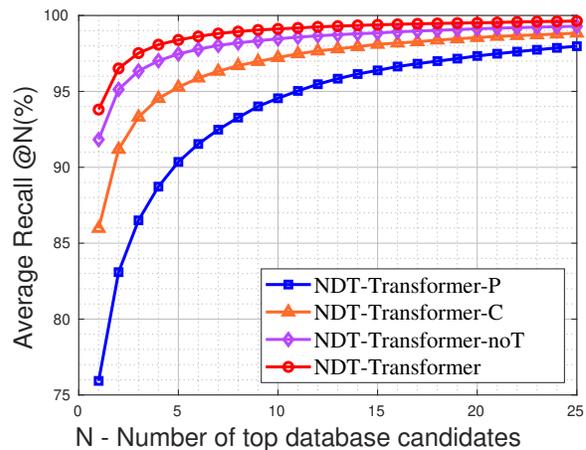


Fig. 6: The average recall curves from top 1 (@1) to top 25 (@25) of the NDT-Transformer ablation studies.

large-scale topological localisation problem by the means of place recognition and retrieval. Our approach leverages the generic NDT encoding to represent the point cloud, where geometrical information is implicitly summarised. Compared to the state-of-the-art method [15], our approach does not require hand-crafted features, and instead aggregates and strengthens local features hierarchically, hence achieving better results with superior run time performance.

The experimental results on the Oxford Robocar dataset demonstrate the effectiveness of our approach, i.e. a significant improvement of 7.52% on average recall of top 1 and 2.73% on top 1% in comparison with the SOTA method. Our system is able to successfully find a compromise between the needs of SOTA performance and run time performance as an essential module of a real-time SLAM system. The proposed approach also provides an important supplement for NDT-based SLAM approaches such as NDT mapping, NDT-MCL, etc.

REFERENCES

- [1] W. Wang, B. Wang, P. Zhao, C. Chen, R. Clark, B. Yang, A. Markham, and N. Trigoni, "Pointloc: Deep pose regressor for lidar point cloud localization," *arXiv preprint arXiv:2003.02392*, 2020.
- [2] L. Sun, D. Adolphsson, M. Magnusson, H. Andreasson, I. Posner, and T. Duckett, "Localising faster: Efficient and precise lidar-based robot localisation in large-scale environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4386–4392.
- [3] N. Akai, T. Hirayama, and H. Murase, "Hybrid localization using model-and learning-based methods: Fusion of monte carlo and e2e localizations via importance sampling," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2020.
- [4] R. Dubé, D. Dugas, E. Stumm, J. Nieto, R. Siegwart, and C. Cadena, "SegMatch: Segment based place recognition in 3d point clouds," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5266–5272.
- [5] R. Dube, A. Cramariuc, D. Dugas, J. Nieto, R. Siegwart, and C. Cadena, "SegMap: 3d segment mapping using data-driven descriptors," in *Robotics: Science and Systems (RSS)*, 2018.
- [6] G. Tinchev, A. Penate-Sanchez, and M. Fallon, "Learning to see the wood for the trees: Deep laser localization in urban and natural environments on a cpu," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1327–1334, 2019.
- [7] 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," *arXiv preprint arXiv:2008.11459*, 2020.
- [8] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [9] L. He, X. Wang, and H. Zhang, "M2dp: A novel 3d point cloud descriptor and its application in loop closure detection," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 231–237.
- [10] G. Kim and A. Kim, "Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4802–4809.
- [11] K. P. Cop, P. V. Borges, and R. Dubé, "Delight: An efficient descriptor for global localisation using lidar intensities," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3653–3660.
- [12] G. Kim, B. Park, and A. Kim, "1-day learning, 1-year localization: Long-term lidar localization using scan context image," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1948–1955, 2019.
- [13] X. Chen, T. Läbe, A. Millioto, T. Röhlting, O. Vysotska, A. Haag, J. Behley, C. Stachniss, and F. Fraunhofer, "Overlapnet: Loop closing for lidar-based slam," in *Proc. of Robotics: Science and Systems (RSS)*, 2020.
- [14] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [15] Z. Liu, S. Zhou, C. Suo, P. Yin, W. Chen, H. Wang, H. Li, and Y.-H. Liu, "Lpd-net: 3d point cloud learning for large-scale place recognition and environment analysis," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 2831–2840.
- [16] M. Angelina Uy and G. Hee Lee, "Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4470–4479.
- [17] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3d classification and segmentation," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 1, no. 2, p. 4, 2017.
- [18] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *Acm Transactions On Graphics (tog)*, vol. 38, no. 5, pp. 1–12, 2019.
- [19] P. Biber and W. Straßer, "The normal distributions transform: A new approach to laser scan matching," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 3. IEEE, 2003, pp. 2743–2748.
- [20] T. Stoyanov, J. Saarinen, H. Andreasson, and A. J. Lilienthal, "Normal distributions transform occupancy map fusion: Simultaneous mapping and tracking in large scale dynamic environments," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4702–4708.
- [21] M. Magnusson, A. Lilienthal, and T. Duckett, "Scan registration for autonomous mining vehicles using 3d-ndt," *Journal of Field Robotics*, vol. 24, no. 10, pp. 803–827, 2007.
- [22] J. Saarinen, H. Andreasson, T. Stoyanov, and A. J. Lilienthal, "Normal distributions transform monte-carlo localization (ndt-mcl)," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 382–389.
- [23] D. Adolphsson, S. Lowry, M. Magnusson, A. Lilienthal, and H. Andreasson, "A submap per perspective-selecting subsets for super mapping that afford superior localization quality," in *2019 European Conference on Mobile Robots (ECMR)*. IEEE, 2019, pp. 1–7.
- [24] M. Magnusson, H. Andreasson, A. Nüchter, and A. J. Lilienthal, "Automatic appearance-based loop detection from three-dimensional laser data using the normal distributions transform," *Journal of Field Robotics*, vol. 26, no. 11-12, pp. 892–914, 2009.
- [25] A. Zaganidis, A. Zernitev, T. Duckett, and G. Cielniak, "Semantically assisted loop closure in slam using ndt histograms," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4562–4568.
- [26] M. Magnusson, T. P. Kucner, S. G. Shahbandi, H. Andreasson, and A. J. Lilienthal, "Semi-supervised 3d place categorisation by descriptor clustering," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 620–625.
- [27] T. Schmiedel, E. Einhorn, and H. Gross, "Iron: A fast interest point descriptor for robust ndt-map matching and its application to robot localization," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 3144–3151.
- [28] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 Year, 1000km: The Oxford RobotCar Dataset," *The International Journal of Robotics Research (IJRR)*, vol. 36, no. 1, pp. 3–15, 2017. [Online]. Available: <http://dx.doi.org/10.1177/0278364916679498>