
DeepSeqSLAM: A Trainable CNN+RNN for Joint Global Description and Sequence-based Place Recognition

Marvin Chancán
QUT, Australia
mchancan1@uni.pe

Michael Milford
QUT, Australia
michael.milford@qut.edu.au

Abstract

Sequence-based place recognition methods for all-weather navigation are well-known for producing state-of-the-art results under challenging day-night or summer-winter transitions. These systems, however, rely on complex handcrafted heuristics for sequential matching—which are applied on top of a pre-computed pairwise similarity matrix between reference and query image sequences of a single route—to further reduce false-positive rates compared to single-frame retrieval methods. As a result, performing multi-frame place recognition can be extremely slow for deployment on autonomous vehicles or evaluation on large datasets, and fail when using relatively short parameter values such as a sequence length of 2 frames. In this paper, we propose DeepSeqSLAM: a trainable CNN+RNN architecture for *jointly learning visual and positional representations* from a single monocular image sequence of a route. We demonstrate our approach on two large benchmark datasets, Nordland and Oxford RobotCar—recorded over 728 km and 10 km routes, respectively, each during 1 year with multiple seasons, weather, and lighting conditions. On Nordland, we compare our method to two state-of-the-art sequence-based methods across the entire route under summer-winter changes using a sequence length of 2 and show that our approach can get over 72% AUC compared to 27% AUC for Delta Descriptors and 2% AUC for SeqSLAM; while drastically reducing the deployment time from around 1 hour to 1 minute against both. The framework code and video are available at [mchancan.github.io/deepseqslam](https://github.com/mchancan/deepseqslam)

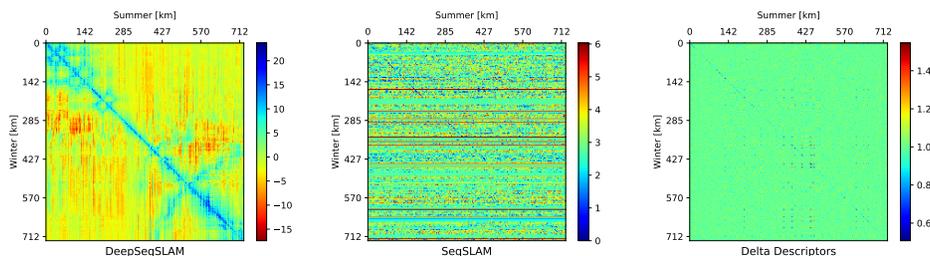


Figure 1: **State-of-the-art comparison of sequential matching between *summer* (reference) and *winter* (query) conditions on the Nordland dataset.** We use a challenging (short) sequence length of 2 frames to evaluate these methods on the full 728 km route of the dataset—comprising 3577 frames per reference/query traversal recorded at 0.1 FPS. **Left:** Neural activity profiles of the linear output layer of DeepSeqSLAM yielding over 72% average precision (higher unit activation (*blue*) represents a strong match candidate). **Middle, Right:** Difference matrices obtained using classical multi-frame methods such as SeqSLAM and Delta Descriptors resulting in 2% and 27% average precision, respectively (lower score (*blue*) means a strong match candidate).

1 Introduction and Related Work

Localization is a key component for enabling the deployment of autonomous vehicles (AV) with full self-driving capabilities in real environments, including a range of other driver assistance features such as place recognition [4, 12, 27, 36, 43], multi-object tracking [49, 14], road scene understanding [5, 24, 23], obstacle detection [22, 29], and also large-scale driving datasets [19, 28, 21, 50, 52] for training autonomous driving (AD) models [51, 1, 2]. In place recognition research, there have been great advances on the use of convolutional neural networks (CNN) [20] for visual place recognition [4, 42, 12, 3, 13] and visual localization tasks under extreme changing conditions [15, 37]. These models have shown to generalize well across challenging appearance and viewpoint variations including multiple seasons, weather, and lighting conditions [12, 3, 37]. On the other hand, multi-frame-based methods for sequential matching are known for improving single-image retrieval results on driving datasets with significant appearance changes [30, 47, 48, 32, 45]. However, state-of-the-art sequence-based approaches like SeqSLAM [30] heavily rely on computationally expensive heuristics to further enhance the difference matrix between query and reference image sequences of the same route. Although researchers have proposed a number of sequence-based methods [35, 18, 17], these often exhibit the fundamental limitations of SeqSLAM: a) computational cost scale linearly with the dataset size, b) fail when using relatively short parameter values such as a sequence length of 2, and c) requiring both reference and query traversals to be pre-aligned with the same number of frames.

In this work, we study the use of recurrent neural networks (RNN) [25] for temporally integrating sequential information over short image sequences [16]; while leveraging the computational and dynamic information processing functions of these networks for end-to-end positional and sequence-based place learning for the first time. Hence, we propose a CNN+RNN architecture that can be trained end-to-end for *jointly learning visual and positional representations*. Our approach also enables the use of any pre-trained CNN model, where the RNN component can then be trained using a single monocular traversal of a driving dataset. In contrast to classical sequence-based methods [30, 35, 18, 17]—which often require velocity information or reference-query image sequences being of the same size—and recent neuro-inspired systems [39, 33, 11], our approach does not require velocity data and can work with query traversals of any size—only requiring positional information of the AV when traversing an environment, which can be obtained from visual odometry (VO) [31], radar odometry (RO) [6], LiDAR, structure from motion (SfM) [38] or any other source of motion estimation. Thus enabling our approach to also work as a full SLAM system by encoding two key sensor modalities for autonomous navigation within a single *motion and visual* perception framework [10]. We also show that our approach is orders of magnitude faster than classical methods and can be deployed on CPU, GPU, or multi-GPUs across many servers.

2 DeepSeqSLAM

In robot navigation research for AD, both visual and positional information share the underlying representation that captures the nature of an environment as the robot navigates its surroundings. Contrary to the classical two-stage place recognition pipelines, i.e. match-then-temporally-filter, we propose to jointly perform *visual and positional representation learning* to obtain bimodal descriptors that can be simultaneously used for performing sequence-based place recognition. We then refer to our approach as DeepSeqSLAM and detail its components in the next Section.

2.1 Global Place Description

Given an image sequence I_t of an environment, we apply a CNN function on each input image to obtain compact n -dimensional global image descriptors \mathbf{d}_t , where $n \in \mathbb{N}^+$ is a function of the CNN model. These representations can be learned through conventional backpropagation in a training stage, and then used to find the top-k similar images across a query sequence of the same route; potentially with different viewpoints and environmental conditions. After training, Cosine or Euclidean distances are often used to compute a similarity matrix between reference and query images.

2.2 Sequence-based Place Learning

Classical multi-frame-based place recognition methods are required to first compute a pairwise difference matrix between both reference and query image sequences to then iteratively apply temporal

Table 1: Two Large-Scale Benchmark Datasets for Sequence-based Place Recognition.

Dataset	Environmental Condition	Journey	Size (frames)	Sensor
Nordland Railway	Summer	728 km	35768	1 × camera (1 FPS)
	Fall	728 km	35768	
	Winter	728 km	35768	
	Spring	728 km	35768	
Oxford RobotCar	Day	10 km	31472	1 × camera (15 FPS)
	Night	10 km	38315	
	Snow	10 km	40943	
	Rain	10 km	37564	

filtering/searching heuristics on top of these results to further reduce false-positive rates. These traditional methods have demonstrated to be highly inefficient and also impractical for deployment on real AVs or large datasets. Particularly when we want to deploy a localization system on a single query traversal after the training stage has been done on a reference route of the same environment.

Here we propose to integrate an RNN model on top of a CNN function for sequence learning from a single traversal of a route. Depending on the CNN model, we can alleviate the large amount of data—typically required for effective training—by plugin a pre-trained CNN-based network such as NetVLAD [4]; heavily used in visual place recognition research for extracting robust global image descriptors. Thus, the CNN function can be freeze while the RNN component is trained. The required input for the RNN is a global image descriptor \mathbf{d}_t (obtained from full-resolution RGB images) concatenated with its corresponding 2D global positional information \mathbf{p}_t (represented here using a 2- d vector). In our experiments, we use the corresponding GPS data normalized between -1 and 1 for \mathbf{p}_t , although other sources of motion estimation such as VO, radar, or LiDAR can be used.

2.3 Implementation Details

We use the NetVLAD [4] network (based on a VGG-16 architecture [40] trained on Pittsburgh 30k [46]) as our back-end CNN function for global image description, with a feature dimension of $n = 4096$ and l_2 -normalized. For the RNN component, a single cell, long short-term memory (LSTM) with 512 units is used, with a multi-layer perceptron (MLP) with N units—equal to the number of frames in the training data—over the top for receiving the output state of the LSTM; which after training is capable of generating the neural activity profiles shown in Fig. 1. We use a learning rate of 0.01 with Adam [26] for training. In sequence-based place recognition, the sequence length, $d_s \in \mathbb{N}^+$, is a sensitive parameter [9] that is typically predefined to arrange the training data in short consecutive overlapping sequences of size d_s frames. We provide extensive experimental results on the trade-off of using different d_s values from 1 to 24, and compare the performance of our method to two state-of-the-art multi-frame systems in the next Section.

3 Experimental Evaluations

3.1 Large-Scale Datasets

In Table 1 we report the details of the datasets used to demonstrate our approach. The Nordland dataset [41] was recorded on a 728 km train journey in Norway, providing four long traversals, once per season, with diverse visual conditions over 1 year. We train our approach on a full traversal (*summer*) and test on the remaining (*fall*, *winter*, *spring*). The Oxford RobotCar dataset [28] was collected on a car platform over a 10 km route in Oxford, UK, also over 1 year. The full dataset includes a range of sensory information from LiDAR, monocular cameras, and trinocular stereo cameras, over +100 traversals with different weather, seasons, and dynamic urban conditions. We select 4 sequences of this dataset¹ and train our method on a single traversal.

3.2 State-of-the-Art Results

In Fig. 2 we study the influence of changing the sequence length on our proposed method trained on the *summer* and compare their generalization capabilities to *fall*, *winter*, and *spring* conditions

¹Ref. as 2015-05-19-14-06-38, 2014-12-10-18-10-50, 2015-02-03-08-45-10, 2015-10-29-12-18-17 in [24].

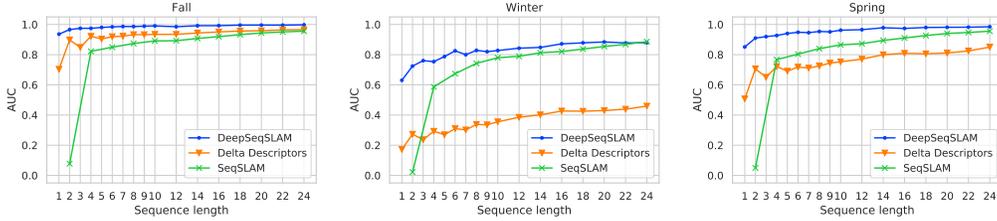


Figure 2: **Influence of sequence length (d_s)** in multi-frame methods and its generalization capabilities from reference (*summer*) to query traversals (*fall*, *winter*, and *spring*) of the Nordland dataset.

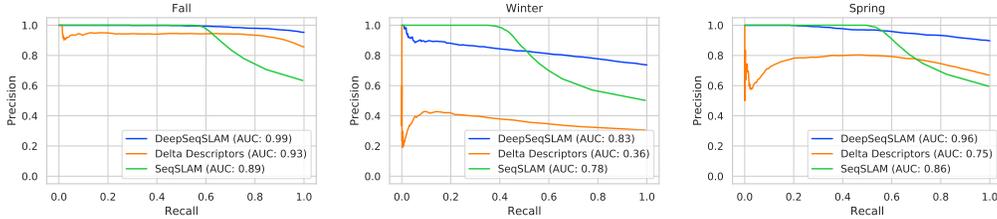


Figure 3: **State-of-the-Art comparison of PR performance on Nordland with sequence length 10.**

with classical sequence-based approaches such as SeqSLAM [30] and the recently published Delta Descriptors work [18]. For SeqSLAM, we use the original MATLAB implementation provided in [44] with default parameters (only changing the sequence length d_s as required); the SeqSLAM algorithm only works using an even number for d_s though. For Delta Descriptors, we use the original Python code provided in [18] with default parameters also varying the sequence length only as needed. We highlight that DeepSeqSLAM is getting such good performance with very short sequence lengths, compared to classical methods. In particular, under *summer* to *winter* changes with a sequence length of 2 (Fig. 2-Middle), SeqSLAM and Delta Descriptors get 2% and 27% AUC, respectively, while our method gets over 72% AUC. In Fig. 1 we provide the full similarity matrix results for each method under this challenging setting. Furthermore, in Fig. 3 we report the PR curves for a sequence length of 10 across all these models, including its corresponding AUC metrics. We note that SeqSLAM performs competitively at this d_s with 78% AUC on *winter*, Fig. 3-middle, but DeepSeqSLAM outperforms it with 83% AUC. For computing the average-precision/AUC results from Figs. 1, 2 and 3, given a *query* image I_{q_t} and its corresponding *reference* image I_{d_t} , the retrieved images between $I_{d_t-\delta}$ and $I_{d_t+\delta}$ were considered as a correct match; where the tolerance δ was given by $d_s + 10$. An overview of the proposed CNN+RNN architecture is presented in Fig. 4 of the **Appendix**, along with its training curves (accuracy and loss traces) in Fig. 5, Table 2 on the computational performance analysis, and qualitative results on the Nordland dataset in Fig. 6.

4 Discussion and Conclusions

The results on the full Nordland dataset show that our proposed CNN+RNN model is capable of learning meaningful temporal relations from a single image sequence of a large driving dataset; while significantly outperforming classical sequence-based methods in runtime, accuracy, and computational requirements. We used a small two-layer CNN for exploring the end-to-end training behavior (from scratch) of DeepSeqSLAM but our preliminary results showed that the CNN component does not generalize well to drastic visual changes; which was expected since these models require a significant amount of data for effectively training and generalizing. We see this observation as future work to further investigate the advantages of jointly learning visual and positional information for AD applications. Finally, we call our method DeepSeqSLAM since it is our goal to build a full simultaneous localization and mapping (SLAM) system using this framework by incorporating a learning-based component for geometric mapping such as those in [8, 7, 53]. Our first attempts to integrate these mapping systems into our framework were limited by different, expensive training requirements from these models, but we found that on the deployment stage our approach can be easily integrated with a mapping system and we are currently exploring this for future work.

References

- [1] Alexander Amini, Guy Rosman, Sertac Karaman, and Daniela Rus. Variational end-to-end navigation and localization. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8958–8964, May 2019.
- [2] Alexander Amini, Igor Gilitschenski, Jacob Phillips, Julia Moseyko, Rohan Banerjee, Sertac Karaman, and Daniela Rus. Learning robust control policies for end-to-end autonomous driving from data-driven simulation. *IEEE Robotics and Automation Letters*, 5(2):1143–1150, April 2020. ISSN 2377-3774. doi: 10.1109/LRA.2020.2966414.
- [3] Mikaela Angelina Uy and Gim 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*, pages 4470–4479, 2018.
- [4] Relja Arandjelović, Petr Gronát, Akihiko Torii, Tomáš Pajdla, and Josef Sivic. NetVLAD: CNN Architecture for Weakly Supervised Place Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5297–5307, 2016.
- [5] Dan Barnes, Will Maddern, and Ingmar Posner. Find your own way: Weakly-supervised segmentation of path proposals for urban autonomy. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 203–210, May 2017.
- [6] Dan Barnes, Matthew Gadd, Paul Murcutt, Paul Newman, and Ingmar Posner. The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6433–6438, 2020.
- [7] Jia-Wang Bian, Zhichao Li, Naiyan Wang, Huangying Zhan, Chunhua Shen, Ming-Ming Cheng, and Ian Reid. Unsupervised scale-consistent depth and ego-motion learning from monocular video. In *Thirty-third Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [8] Samarth Brahmabhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-aware learning of maps for camera localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [9] Jack Bruce, Adam Jacobson, and Michael Milford. Look no further: Adapting the localization sensory window to the temporal characteristics of the environment. *IEEE Robotics and Automation Letters*, 2(4): 2209–2216, 2017.
- [10] Marvin Chancán and Michael Milford. Mvp: Unified motion and visual self-supervised learning for large-scale robotic navigation. *arXiv preprint arXiv:2003.00667*, 2020.
- [11] Marvin Chancán, Luis Hernandez-Nunez, Ajay Narendra, Andrew B. Barron, and Michael Milford. A hybrid compact neural architecture for visual place recognition. *IEEE Robotics and Automation Letters*, 5(2):993–1000, April 2020. ISSN 2377-3774. doi: 10.1109/LRA.2020.2967324.
- [12] Zetao Chen, Adam Jacobson, Niko Sünderhauf, Ben Upcroft, Lingqiao Liu, Chunhua Shen, Ian D. Reid, and Michael Milford. Deep learning features at scale for visual place recognition. *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3223–3230, 2017.
- [13] Zetao Chen, Lingqiao Liu, Inkyu Sa, Zongyuan Ge, and Margarita Chli. Learning context flexible attention model for long-term visual place recognition. *IEEE Robotics and Automation Letters*, 3(4):4015–4022, 2018.
- [14] Hsu-kuang Chiu, Antonio Prioletti, Jie Li, and Jeannette Bohg. Probabilistic 3d multi-object tracking for autonomous driving. *arXiv preprint arXiv:2001.05673*, 2020.
- [15] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint description and detection of local features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8092–8101, 2019.
- [16] Jose M Facil, Daniel Olid, Luis Montesano, and Javier Civera. Condition-invariant multi-view place recognition. *arXiv preprint arXiv:1902.09516*, 2019.
- [17] Matthew Gadd, Daniele De Martini, and Paul Newman. Look around you: Sequence-based radar place recognition with learned rotational invariance. In *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 270–276, 2020.

- [18] Sourav Garg, Ben Harwood, Gaurangi Anand, and Michael Milford. Delta descriptors: Change-based place representation for robust visual localization. *IEEE Robotics and Automation Letters*, 5(4):5120–5127, 2020.
- [19] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, June 2012. doi: 10.1109/CVPR.2012.6248074.
- [20] Ian G Goodfellow, Yoshua Bengio, and Aaron C. Courville. Deep learning. *Nature*, 521:436–444, 2015.
- [21] Junyao Guo, Unmesh Kurup, and Mohak Shah. Is it safe to drive? an overview of factors, challenges, and datasets for driveability assessment in autonomous driving. *arXiv preprint arXiv:1811.11277*, 2018.
- [22] Christian Häne, Torsten Sattler, and Marc Pollefeys. Obstacle detection for self-driving cars using only monocular cameras and wheel odometry. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5101–5108. IEEE, 2015.
- [23] Christian Häne, Lionel Heng, Gim Hee Lee, Friedrich Fraundorfer, Paul Furgale, Torsten Sattler, and Marc Pollefeys. 3d visual perception for self-driving cars using a multi-camera system: Calibration, mapping, localization, and obstacle detection. *Image and Vision Computing*, 68:14–27, 2017.
- [24] Lionel Heng, Benjamin Choi, Zhaopeng Cui, Marcel Geppert, Sixing Hu, Benson Kuan, Peidong Liu, Rang Nguyen, Ye Chuan Yeo, Andreas Geiger, et al. Project autovision: Localization and 3d scene perception for an autonomous vehicle with a multi-camera system. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4695–4702, 2019.
- [25] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9:1735–1780, 1997.
- [26] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [27] Stephanie Lowry, Niko Sünderhauf, Paul Newman, John J Leonard, David Cox, Peter Corke, and Michael J Milford. Visual place recognition: A survey. *IEEE Transactions on Robotics*, 32(1):1–19, Feb 2016.
- [28] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017.
- [29] Rowan McAllister, Yarin Gal, Alex Kendall, Mark van der Wilk, Amar Shah, Roberto Cipolla, and Adrian Weller. Concrete problems for autonomous vehicle safety: Advantages of bayesian deep learning. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 4745–4753, 2017.
- [30] Michael J Milford and Gordon F Wyeth. SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights. In *2012 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1643–1649, May 2012.
- [31] Sherif AS Mohamed, Mohammad-Hashem Haghbayan, Tomi Westerlund, Jukka Heikkonen, Hannu Tenhunen, and Juha Plosila. A survey on odometry for autonomous navigation systems. *IEEE Access*, 7: 97466–97486, 2019.
- [32] Tayyab Naseer, Wolfram Burgard, and Cyrill Stachniss. Robust visual localization across seasons. *IEEE Transactions on Robotics*, 34(2):289–302, 2018.
- [33] Peer Neubert, Stefan Schubert, and Peter Protzel. A neurologically inspired sequence processing model for mobile robot place recognition. *IEEE Robotics and Automation Letters*, 4(4):3200–3207, 2019.
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8026–8037. Curran Associates, Inc., 2019.
- [35] Edward Pepperell, Peter I Corke, and Michael J Milford. All-environment visual place recognition with smart. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1612–1618, 2014.
- [36] Torsten Sattler, Michal Havlena, Konrad Schindler, and Marc Pollefeys. Large-scale location recognition and the geometric burstiness problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1582–1590, 2016.

- [37] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, et al. Benchmarking 6dof outdoor visual localization in changing conditions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8601–8610, 2018.
- [38] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [39] Stefan Schubert, Peer Neubert, and Peter Protzel. Towards combining a neocortex model with entorhinal grid cells for mobile robot localization. In *2019 European Conference on Mobile Robots (ECMR)*, pages 1–8, 2019.
- [40] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [41] Niko Sünderhauf, Peer Neubert, and Peter Protzel. Are we there yet? challenging seqslam on a 3000 km journey across all four seasons. In *Proc. Workshop Long-Term Autonomy 2013 IEEE Int. Conf. Robot. Autom. (ICRA)*, 2013.
- [42] Niko Sünderhauf, Sareh Shirazi, Feras Dayoub, Ben Upcroft, and Michael Milford. On the performance of convnet features for place recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4297–4304, 2015.
- [43] Niko Sünderhauf, Sareh Shirazi, Adam Jacobson, Feras Dayoub, Edward Pepperell, Ben Upcroft, and Michael Milford. Place recognition with convnet landmarks: Viewpoint-robust, condition-robust, training-free. *Proceedings of Robotics: Science and Systems XII*, 2015.
- [44] Ben Talbot, Sourav Garg, and Michael Milford. OpenSeqSLAM2.0: An Open Source Toolbox for Visual Place Recognition Under Changing Conditions. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7758–7765, Oct 2018. doi: 10.1109/IROS.2018.8593761.
- [45] Carl Toft, Will Maddern, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, et al. Long-term visual localization revisited. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [46] Akihiko Torii, Josef Sivic, Tomas Pajdla, and Masatoshi Okutomi. Visual place recognition with repetitive structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(11):2346–2359, 2015.
- [47] Olga Vysotska and Cyrill Stachniss. Lazy data association for image sequences matching under substantial appearance changes. *IEEE Robotics and Automation Letters*, 1(1):213–220, 2015.
- [48] Olga Vysotska, Tayyab Naseer, Luciano Spinello, Wolfram Burgard, and Cyrill Stachniss. Efficient and effective matching of image sequences under substantial appearance changes exploiting gps priors. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2774–2779. IEEE, 2015.
- [49] Dequan Wang, Coline Devin, Qi-Zhi Cai, Fisher Yu, and Trevor Darrell. Deep object-centric policies for autonomous driving. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8853–8859, May 2019.
- [50] Frederik Warburg, Soren Hauberg, Manuel López-Antequera, Pau Gargallo, Yubin Kuang, and Javier Civera. Mapillary street-level sequences: A dataset for lifelong place recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2626–2635, 2020.
- [51] Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell. End-to-end learning of driving models from large-scale video datasets. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [52] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2636–2645, 2020.
- [53] Wang Zhao, Shaohui Liu, Yezhi Shu, and Yong-Jin Liu. Towards better generalization: Joint depth-pose learning without posenet. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

A DeepSeqSLAM architecture

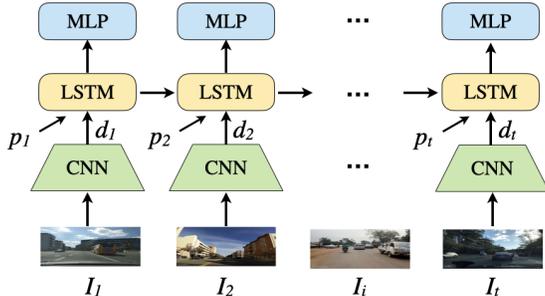


Figure 4: Proposed CNN+RNN architecture for learning positional p_i and visual d_i representations.

B Training curves

In Fig. 5 we present the training curves of our proposed approach on a full traversal of the Nordland and Oxford dataset; *summer* and *day*, respectively. Our results and state-of-the-art comparison on the Oxford dataset are currently in preparation, although the SeqSLAM algorithm does not achieve state-of-the-art performance on this dataset due to significant viewpoint changes and dynamic objects. While in our system, the CNN component makes our method robust to viewpoint variations.

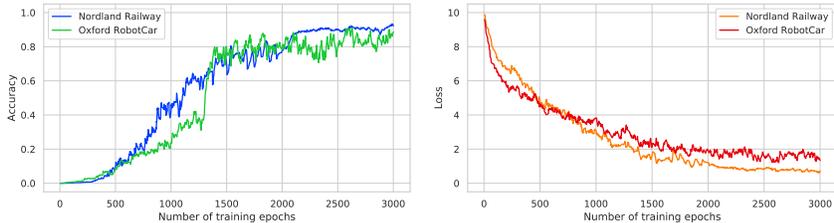


Figure 5: Training curves of DeepSeqSLAM on a single entire traversal of the Nordland Railway and the Oxford RobotCar datasets.

C Computational Performance Analysis

In contrast to traditional methods, that can only run on CPU and are computationally expensive on large datasets, our PyTorch [34] implementation of DeepSeqSLAM has been designed to run on CPU, GPU, and multi-GPUs across many servers. We plan to evaluate DeepSeqSLAM on much larger datasets and also use it for training novel CNN architectures that may require preprocessing stages such as image augmentation for instance. The full source-code of DeepSeqSLAM is made publicly available for the benefit of the community.

Table 2: Computational efficiency on deployment

Method	Runtime on a full query traversal of the Nordland dataset		
	Sequential Matching	Size (frames)	Device
DeepSeqSLAM*	1 min	35768	CPU
SeqSLAM	70 min	35768	CPU
Delta Descriptors	51 min	35768	CPU

*DeepSeqSLAM can also run on a single GPU or even multiple GPUs across many servers.

D Qualitative Results on the Nordland Railway Dataset



Figure 6: **Qualitative results on the full 728km route of the Nordland Railway dataset.** The reference and query image sequences are *summer* and *winter*, respectively, for each method.