IROS '24 ABU DHABI

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

Moving Object Segmentation in Point Cloud Data using Hidden Markov Models

Vedant Bhandari, Jasmin James, Tyson Phillips and P. Ross McAree

The School of Mechanical and Mining Engineering, The University of Queensland

Long-Term Perception for Autonomy in Dynamic Human-centric Environments: What Do Robots Need?

The Moving Object Segmentation Problem

Detecting motion in the agent's workspace is a crucial capability for making informed decisions.

Given a sequence of scans and corresponding sensor pose, the objective is to provide pointwise dynamic classification.

We identify gaps in,

- demonstrating accurate generalized dynamic detection among sensor characteristics and environments, and
- providing a minimal configuration and easy-to-use application.

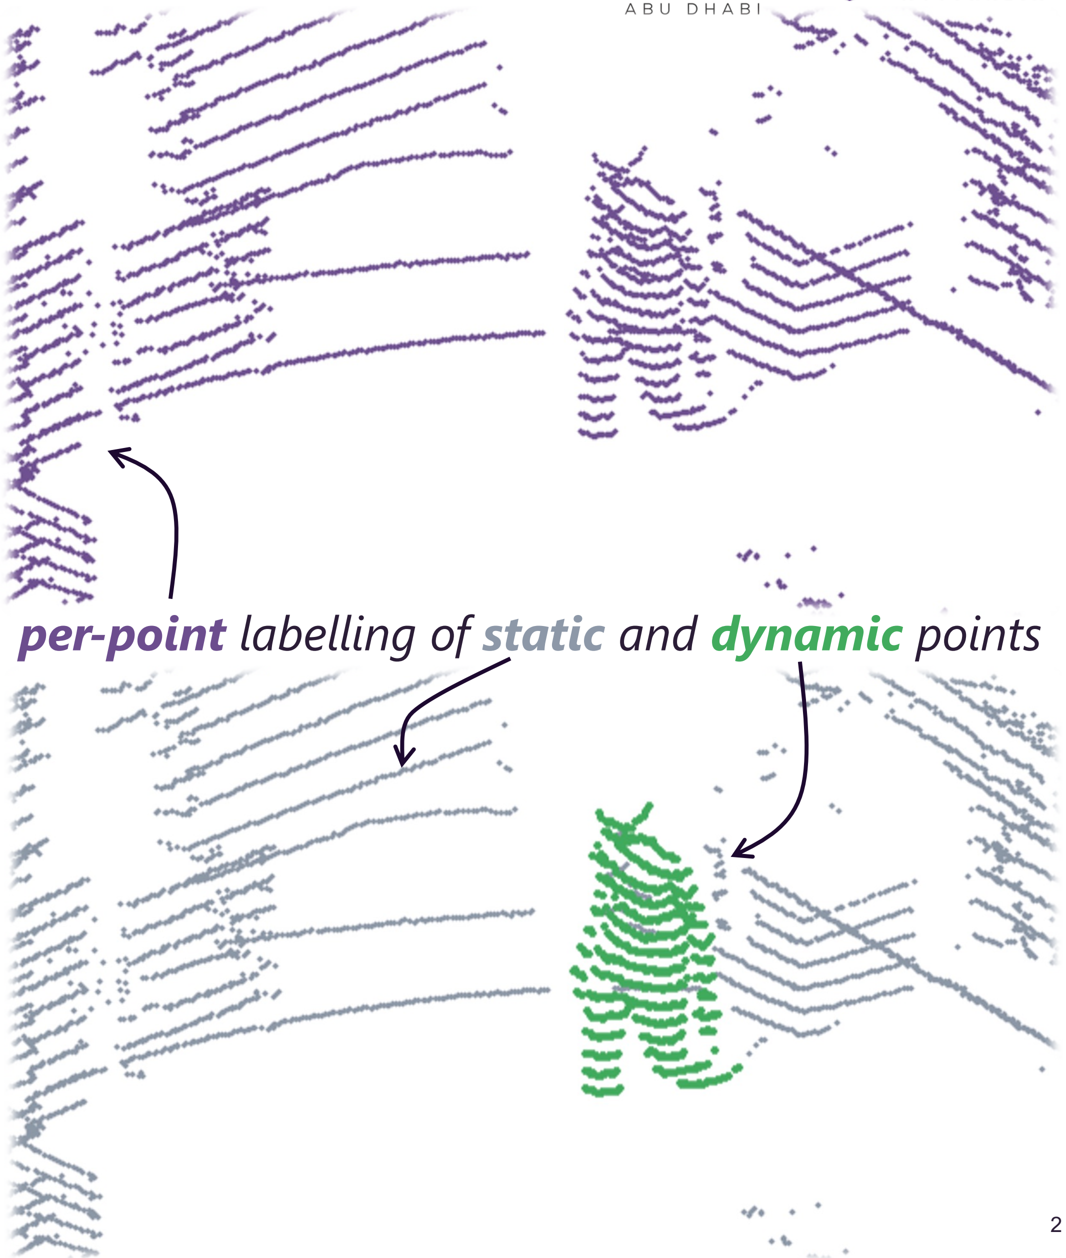

Voxel Representation using HMMs

- We propose a novel learning-free approach to segment moving objects in point cloud data.
- The foundation of the approach lies in modelling each voxel using a hidden Markov model (HMM).

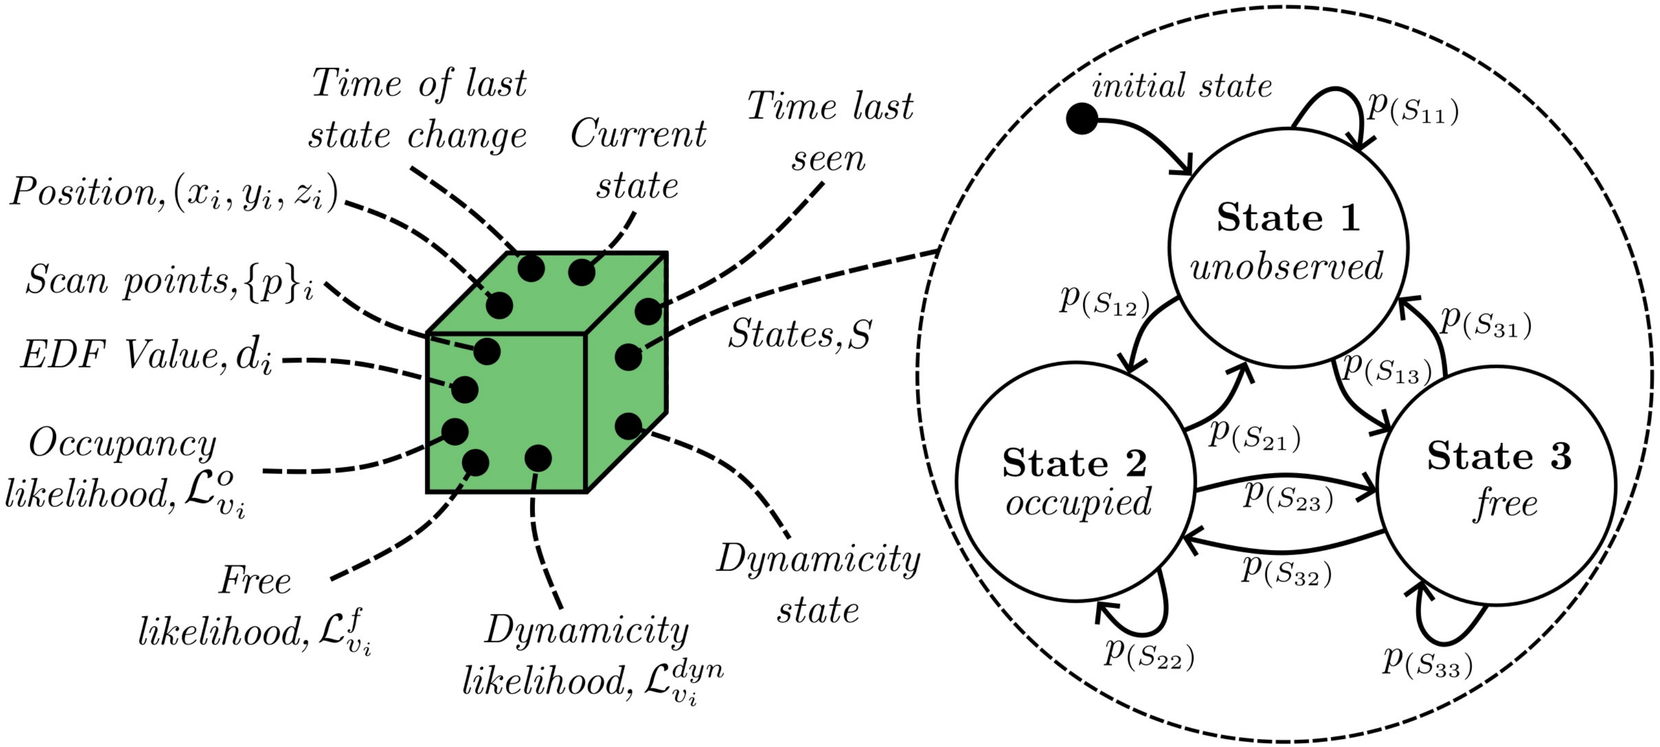

- Each voxel is represented using an HMM with several attributes to encode temporal properties.

The proposed approach uses a simple low-configuration three-stage process to identify dynamic points in a scan.

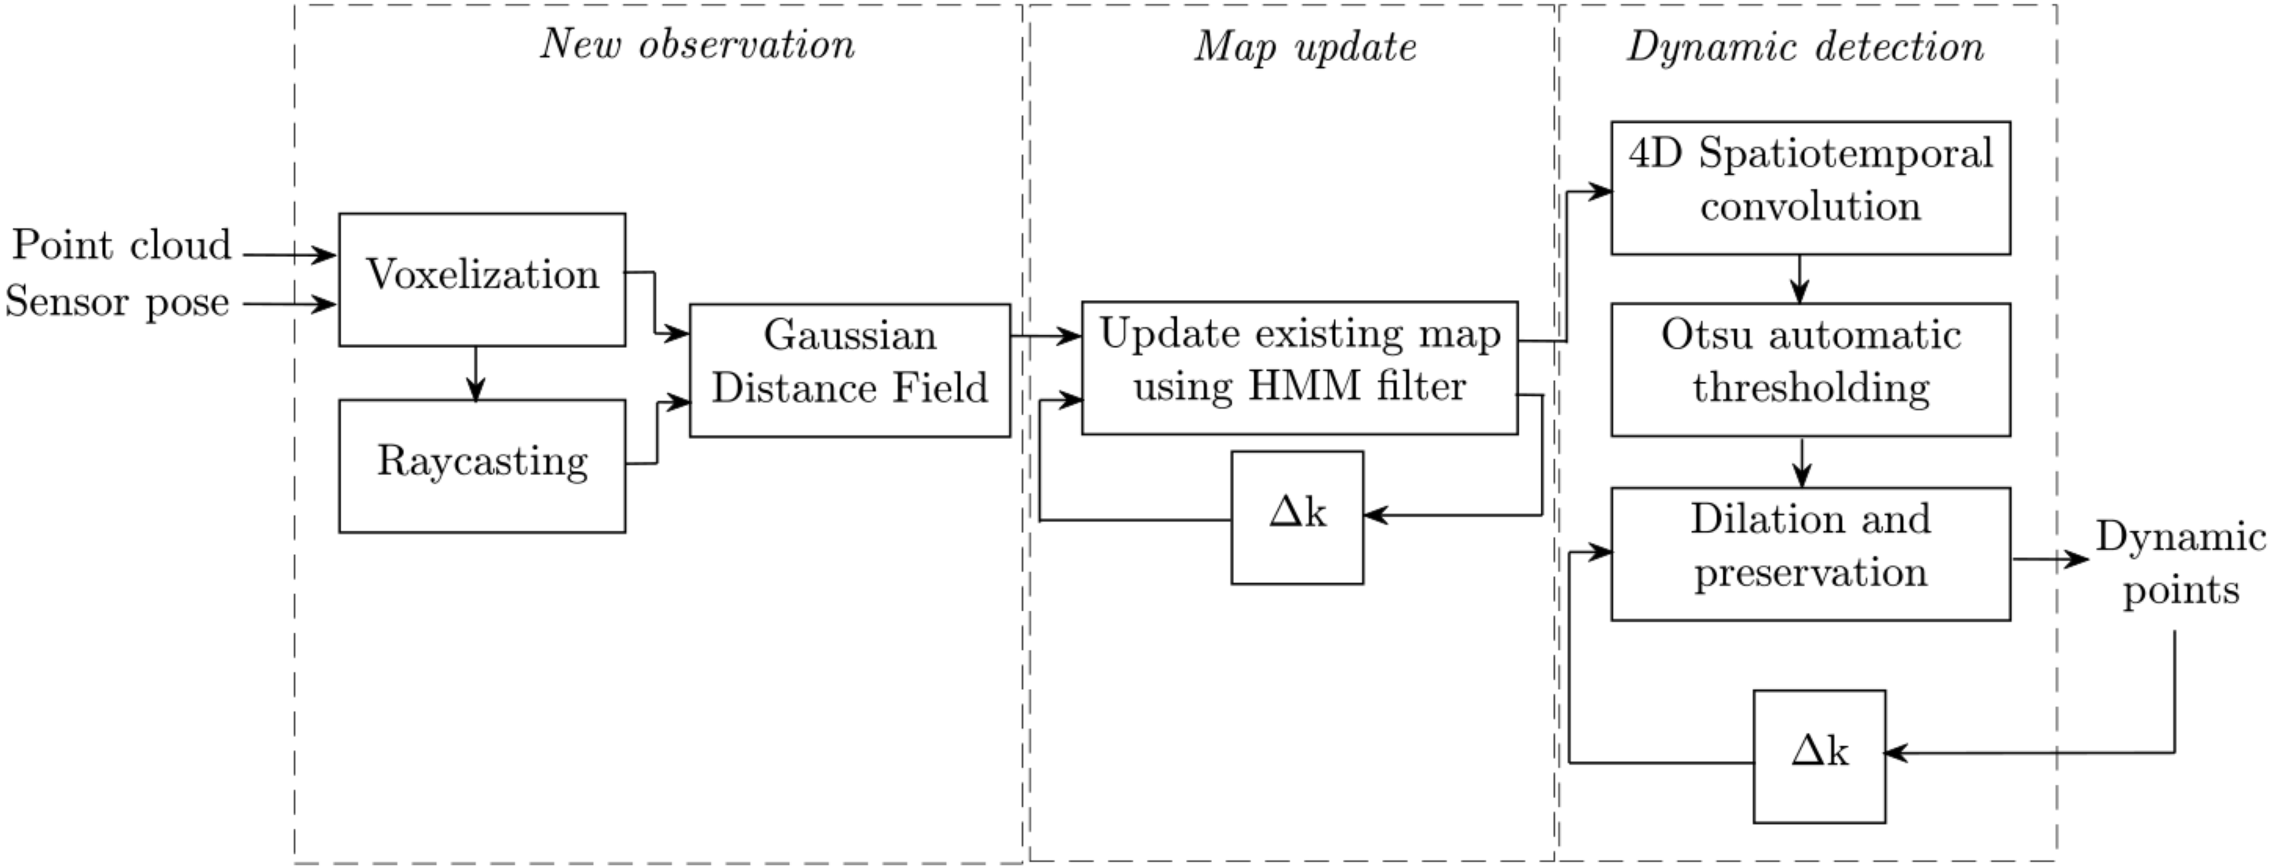

Using 4D Convolutions to filter changes

- 4D Spatiotemporal convolutions help improve true detections (recall) while minimising the false positives.
- Using automatic Otsu thresholding allows for binary class separation for different sensor characteristics, object dynamics, and environments.

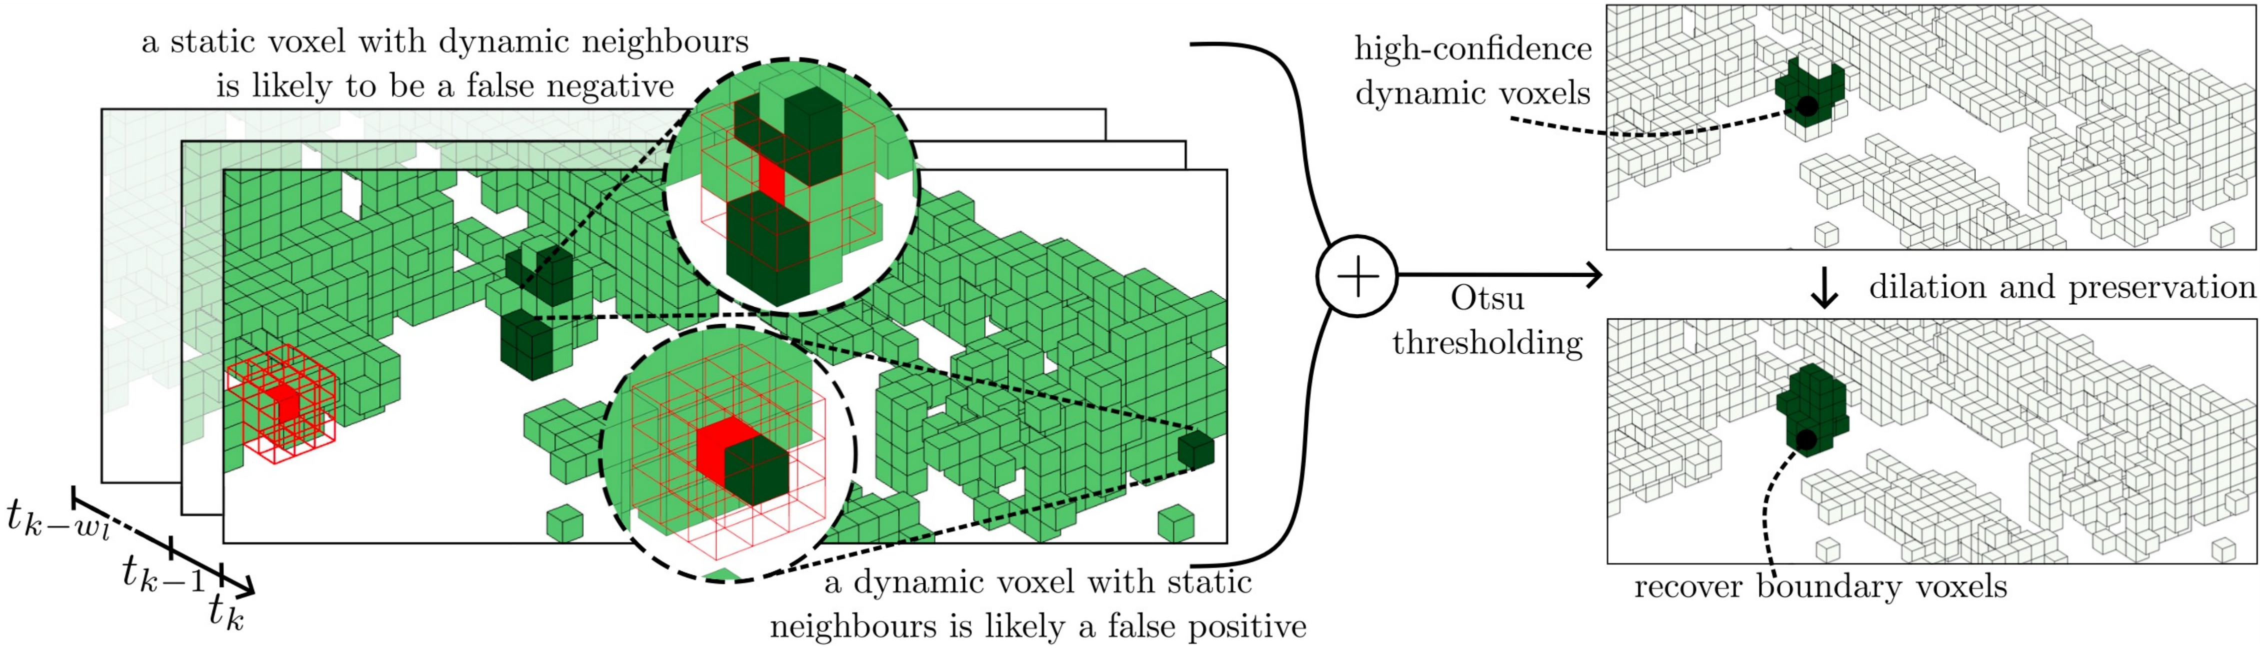

Benchmarking Performance

The proposed approach is benchmarked on numerous datasets: HeLiMOS[1], DOALS[2], Sipailou Campus[3], Apollo[4], Dynablox[5]. Results available on the open-source page[6]!

We achieve consistent accurate performance with the same configuration parameters across all scenarios.

We benchmark performance for detecting objects currently in motion with the option to include a temporal history of dynamic objects.

- Objects currently moving.
- Objects that were moving and are now static.
- Objects that are currently static but move at a future time.
- Objects that have the potential to move but remain static.

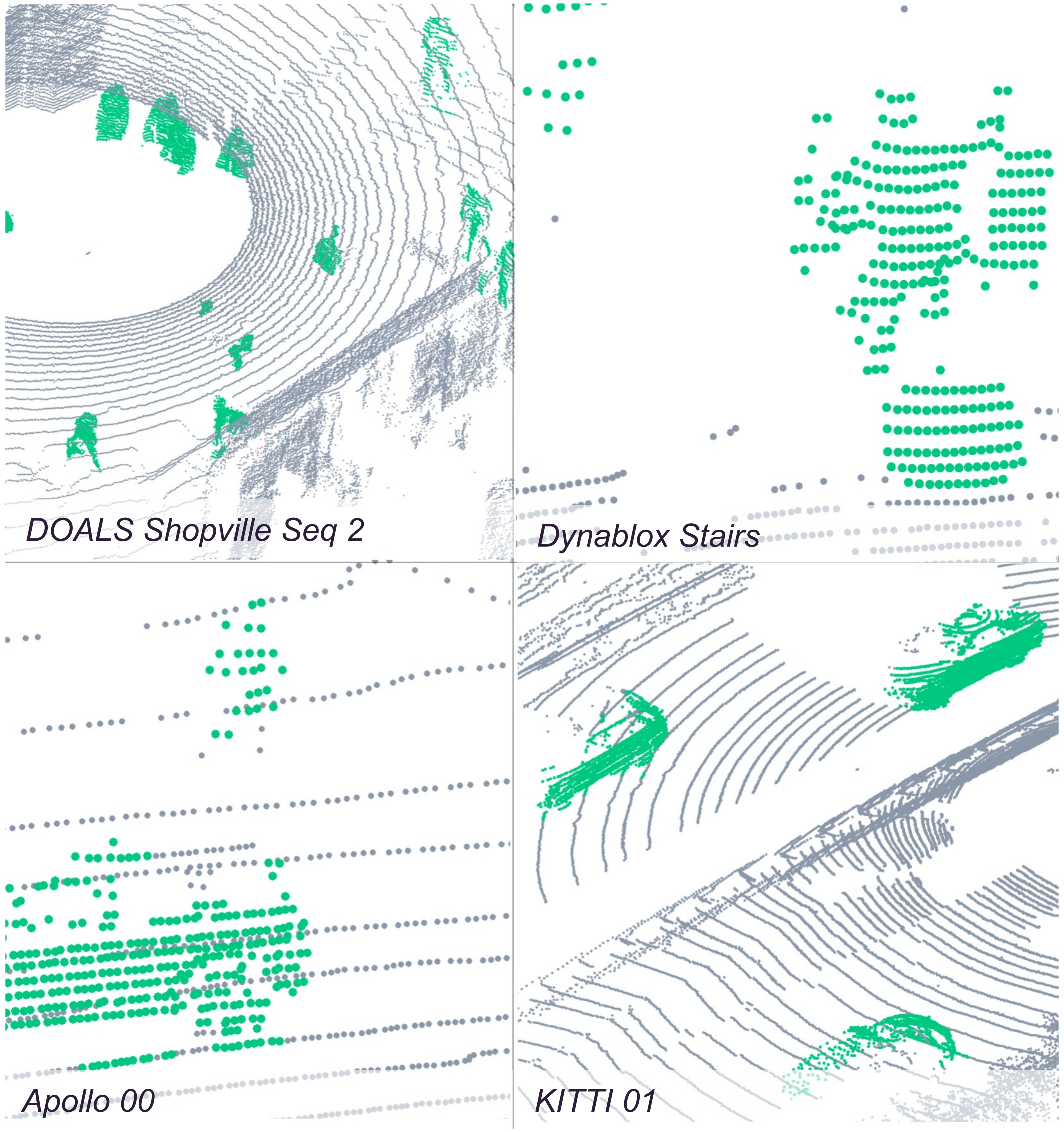

	IoU (%)	Sequences			
	Method	**ST**	**SV**	**HG**	**ND**
DOALS	DOALS-3DMiniNet[2]	84.0	82.0	82.0	80.0
	4DMOS[7]	38.8	50.6	71.1	40.2
	LMNet[8]	19.9	18.9	27.4	40.1
	Dynablox[5]	**86.2**	**83.2**	84.1	**81.6**
	Proposed Approach	82.7	80.8	**85.9**	81.4
	LC Free Space[9] (20 m)	48.7	31.9	24.7	17.7
	Dynablox[5] (20 m)	87.3	**87.8**	86.0	83.1
	Proposed Approach (20 m)	**88.9**	84.7	**87.3**	**83.5**

	IoU (%)	Solid state		Omnidirectional	
	Method	**Livox**	**Aeva**	**OS-128**	**VLP-16**
HeLiMOS	4DMOS[8], online	52.1	54.0	64.2	4.7
	4DMOS[8], delayed	59.0	58.3	70.4	5.4
	MapMOS[10], Scan	58.9	63.2	81.4	4.3
	MapMOS[10], Volume	**62.7**	66.6	**82.9**	5.8
	Proposed Approach	51.3	69.8	75.0	35.0
	Proposed Approach, delayed	57.6	**70.0**	73.4	**53.9**

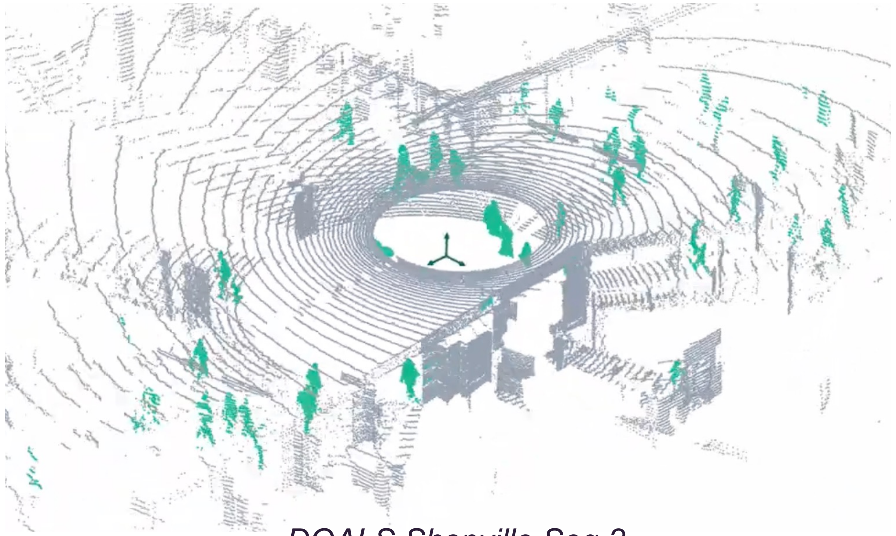

DOALS Shopville Seq 2

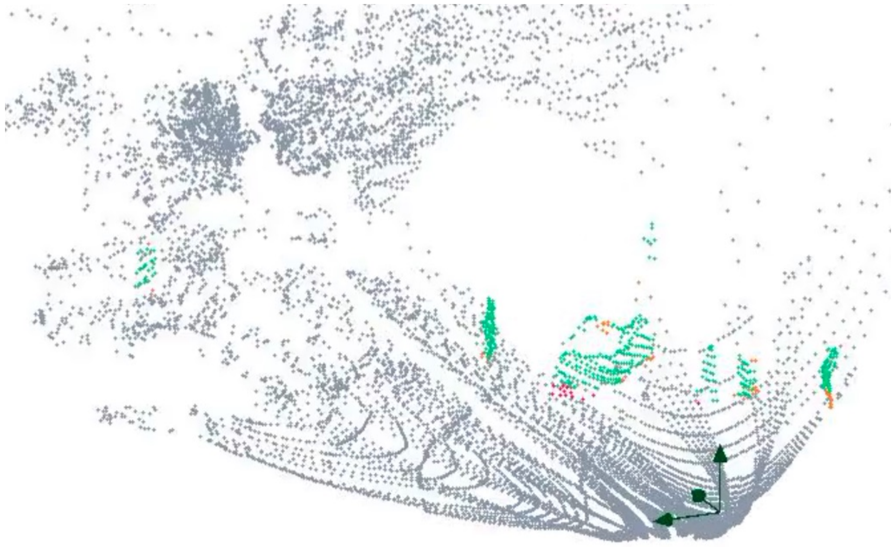

HeLiMOS Avia

Limitations and Future Work

- The current implementation only provides real-time results for 20-50m ranges depending on sensor sparsity.
- While not sensitive, configuring the window size for dynamic memory depends on the situation. Is there a more principled approach to transitioning between dynamic and static states?
- We can currently detect dynamic points. It would be beneficial to provide a means to reflect the varying dynamicity of objects with more meaningful labels.
  - *What does the robot's need?*
  - *Can we provide that information via simple configuration changes?*
- We want the ability to provide an informed belief of the situation to develop a rich decision space.

References

[1] H. Lim, S. Jang, B. Mersch, J. Behley, H. Myung, and C. Stachniss, “Helimos: A dataset for moving object segmentation in 3d point clouds from heterogeneous lidar sensors,” arXiv preprint arXiv:2408.06328, 2024.

[2] P. Pfreundschuh, H. F. Hendrikx, V. Reijgwart, R. Dube, R. Siegwart, ´ and A. Cramariuc, “Dynamic object aware lidar slam based on automatic generation of training data,” in 2021 IEEE International Conference on Robotics and Automation (ICRA), 2021, pp. 11 641– 11 647.

[3] B. Zhou, J. Xie, Y. Pan, J. Wu, and C. Lu, “Motionbev: Attentionaware online lidar moving object segmentation with bird’s eye view based appearance and motion features,” IEEE Robotics and Automation Letters, vol. 8, no. 12, pp. 8074–8081, 2023.

[4] W. Lu, Y. Zhou, G. Wan, S. Hou, and S. Song. L3-Net: Towards Learning Based LiDAR Localization for Autonomous Driving. In Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR), 2019.

[5] L. Schmid, O. Andersson, A. Sulser, P. Pfreundschuh, and R. Siegwart, “Dynablox: Real-time detection of diverse dynamic objects in complex environments,” IEEE Robotics and Automation Letters, vol. 8, no. 10, pp. 6259–6266, 2023.

[6] V. Bhandari, J. James, T. Phillips and P. R. McAree, “Moving Object Segmentation in Point Cloud Data using Hidden Markov Models,” available at https://github.com/vb44/HMM-MOS, 2024.

[7] B. Mersch, X. Chen, I. Vizzo, L. Nunes, J. Behley, and C. Stachniss, “Receding moving object segmentation in 3d lidar data using sparse 4d convolutions,” IEEE Robotics and Automation Letters, vol. 7, no. 3, pp. 7503–7510, 2022.

[8] X. Chen, S. Li, B. Mersch, L. Wiesmann, J. Gall, J. Behley, and C. Stachniss, “Moving object segmentation in 3d lidar data: A learning-based approach exploiting sequential data,” IEEE Robotics and Automation Letters, vol. 6, no. 4, pp. 6529–6536, 2021.

[9] J. Modayil and B. Kuipers, “The initial development of object knowledge by a learning robot,” Robotics and Autonomous Systems, vol. 56, no. 11, pp. 879–890, 2008.

[10] B. Mersch, T. Guadagnino, X. Chen, I. Vizzo, J. Behley, and C. Stachniss, “Building volumetric beliefs for dynamic environments exploiting map-based moving object segmentation,” IEEE Robotics and Automation Letters, vol. 8, no. 8, pp. 5180–5187, 2023.