Semantic Segmentation of Cargo Tank Point Clouds Using Structural Priors

Nikhil Khedekar and Kostas Alexis

***Abstract*—Efficient autonomous inspection of cargo tanks in maritime vessels requires semantic understanding of the tank's structural components to target corrosion-prone members. To this end, we present a geometry-driven pipeline for semantic segmentation of 3D point clouds acquired inside cargo tanks. The method uses the known rectangular prismatic structure of mid-section cargo tanks to estimate a tank aligned coordinate frame from surface normal statistics, fit hull surfaces and the hopper plate via guided Random Sample Consensus (RANSAC), and detect webframes and stringer decks through histogram-based peak detection followed by guided RANSAC plane fitting with density-based clustering. The approach determines the number of webframes and stringer decks without prior knowledge of their count or positions, and operates on noisy, accumulated Light Detection and Ranging (LiDAR) maps. We evaluate on point clouds from multiple crude oil tankers collected by a commercial aerial inspection robot, detecting all seven hull surfaces, five webframes, and four stringer decks, with webframe spacings matching CAD reference values to within 0.01 m across all four tested tanks.**

***Index Terms*—Pointcloud Semantic Segmentation, Maritime Inspection**

I. Introduction

Autonomous robotic inspection of maritime cargo tanks requires the robot to identify specific structural members for detailed examination. We focus on cargo oil tankers, whose mid-section tanks approximate rectangular prisms bounded by flat steel plates: the *deck* (ceiling), *bottom*, two *side shells* (port and starboard walls), and two *bulkheads* (fore and aft end walls). Inside the tank, *webframes* are vertical transverse plates that stiffen the hull, *stringer decks* are horizontal platforms at fixed heights along the walls, *hopper plates* are angled plates connecting the side shell to the bottom, and *longitudinals* are horizontal stiffeners running the length of the tank. Classification societies require detailed examination of these structural members, as they are the primary indicators of structural integrity and corrosion susceptibility [1], [2]. The ability to automatically identify these components enables targeted inspection, where the robot can prioritize high-value regions such as webframe connections and hopper plates.

Training data for tank interiors is scarce as collecting it requires taking the vessel out of normal operation, each tank type has a different geometry, and the point clouds can be noisy. Prior work on autonomous vessel tank inspection has demonstrated aerial robots in ballast water tanks [1], segmented longitudinals in post-processing [3], and detected structural elements geometrically for localization [4], but none address automated semantic segmentation of cargo tank interiors. Learning-based approaches have been explored, including domain adaptation with synthetic data [5], probabilistic deep segmentation for corrosion detection [6], [7], and underwater hull inspection benchmarks [8]. A detailed review of these can be found in [2]. However, these methods focus on image-based inspection rather than 3D structural understanding and require labeled data that is difficult to obtain for cargo tank interiors. Outside the maritime domain, geometry-driven point cloud segmentation works well in structured environments: Li and Shan [9] use constrained Random Sample Consensus (RANSAC) for building reconstruction and Tang et al. [10] combine 2D–3D geometric constraints for indoor space segmentation, though these target architectural rather than maritime structures.

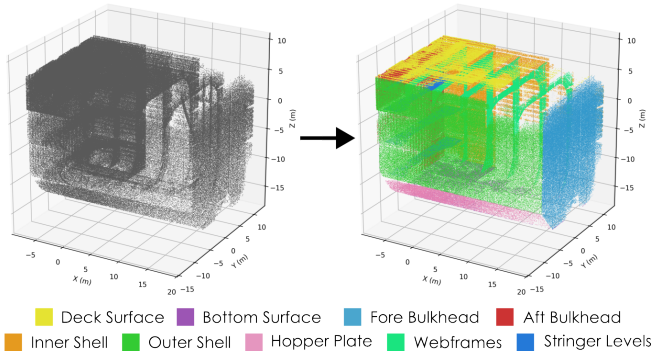

Fig. 1. Left: Accumulated LiDAR point cloud of a starboard side cargo tank collected by an aerial robot. Right: Semantically segmented structural components - hull surfaces, hopper plate, webframes and stringer levels.

We contribute a geometry-driven approach to semantic segmentation of cargo tank point clouds that exploits the structural prior knowledge of the tank and hence requires no training data. We assume: (i) the point cloud is gravity-aligned (from LiDAR-Inertial odometry), (ii) the tank is approximately a rectangular prism with the longitudinal axis longer than the transverse axis (as in the mid-section of a cargo oil tanker), and (iii) the tank side (port or starboard) is known.

II. Approach

We use LiDAR-inertial odometry [11] to produce a gravity-aligned, accumulated point cloud of the tank interior. Each detection stage first isolates the subset of points most likely to

This work was supported by the Research Council of Norway under the AiDiT project (grant number 346524).

The authors are with the Department of Engineering Cybernetics, O. S. Bragstads Plass 2D, Norwegian University of Science and Technology (NTNU), Trondheim, Norway. nikhil.v.khedekar@ntnu.no

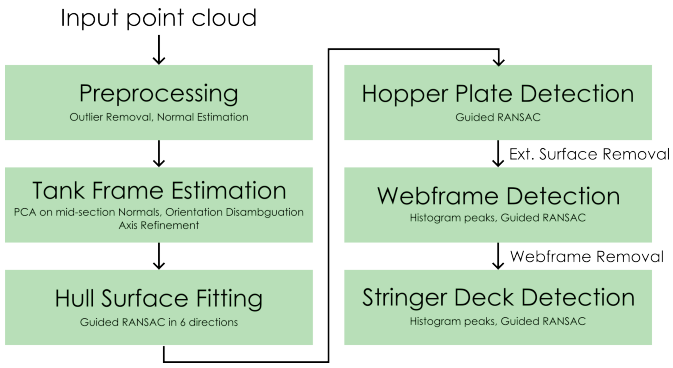

Fig. 2. Overview of the proposed pipeline. The input point cloud is preprocessed and used to estimate the tank coordinate frame. Hull surfaces and the hopper plate are then fitted via guided RANSAC; webframes and stringer decks are subsequently detected on the remaining internals cloud via histogram-based peak detection and guided RANSAC.

belong to the target feature through reduced-space selection, normal alignment filtering, and extremity thresholding. It then applies RANSAC plane fitting to that constrained subset. Because the candidate set is already dominated by the target surface, each fit converges with few iterations. All RANSAC fits use an inlier distance threshold of 0.1 m.

An overview of our pipeline is shown in Fig. 2. Our pipeline proceeds in six stages: preprocessing, tank frame estimation with orientation disambiguation, hull surface fitting, hopper plate detection, webframe detection, and stringer deck detection.

A. Preprocessing

Because hull detection relies on finding extremal points, isolated noise must be removed first. We discard invalid points and apply statistical outlier removal ($k = 50$, 3σ threshold). Surface normals are estimated with the viewpoint at the cloud centroid so that normals point inward; points with invalid normals are discarded.

B. Tank Coordinate Frame Estimation

1) Principal Component Analysis (PCA)-based Axis Estimation: The gravity-alignment assumption fixes the vertical axis to $\mathbf{e}_z = (0, 0, 1)$. To determine the longitudinal and transverse axes, we extract a horizontal slice of points within ± 0.5 m of the centroid height. The normal covariance matrix $\mathbf{C} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{n}_i \mathbf{n}_i^\top$ (where $\mathbf{n}_i$ is the unit-length surface normal for point i) has three eigenvectors. Since the tank walls dominate the normals in the horizontal slice, the two eigenvectors with the largest eigenvalues correspond to the horizontal tank axes. The longitudinal axis is identified as the one along which the point cloud has a greater spatial extent, since cargo tanks are elongated along the ship's length. At the end of this step, we have a preliminary tank frame located at the centroid of the pointcloud with $\mathbf{e}_z$ vertical, $\mathbf{e}_x$ along the ship's length (longitudinal), and $\mathbf{e}_y$ across the ship's width (transverse).

2) Orientation disambiguation: The PCA-based axis estimation leaves a 180° ambiguity in both the longitudinal and transverse directions. We resolve this by detecting the hopper plate. Since the hopper plate always lies on a known side of the tank (determined by whether the tank is port or starboard), its detection uniquely determines the correct axis orientation. We transform the cloud into the preliminary tank frame, extract the likely candidates of the hopper plate as described in Sec. II-D and attempt to fit a plane to it using RANSAC. If the hopper plate is detected on the opposite side from expected, both the longitudinal and transverse axes are flipped. Note that this is a preliminary fit used only for disambiguation; the hopper plate is re-fitted in the refined tank frame in Sec. II-D for the final segmentation result.

3) Axis refinement: We refine the transverse axis by fitting a RANSAC plane to the outer side shell. Candidate points are those whose normals align with the estimated outward-facing direction and that lie within 1 m of the transverse extremity, thereby excluding internal structures parallel to the hull. The fitted plane normal gives the refined transverse axis, and the longitudinal axis follows from $\mathbf{e}_x = \mathbf{e}_y \times \mathbf{e}_z$ (right-hand rule). The cloud is then transformed into this tank-aligned frame.

C. Hull Surface Fitting

With the cloud in the tank frame, we fit the six cardinal hull surfaces (bottom, deck, inner side shell, outer side shell, fore bulkhead, aft bulkhead) using guided RANSAC. The seventh surface, the hopper plate, is handled separately in Sec. II-D. For each of the six cardinal directions $\{\pm\mathbf{e}_x, \pm\mathbf{e}_y, \pm\mathbf{e}_z\}$ in the tank frame, we:

1) Select points in the corresponding half-space whose normals are within 45° of the expected hull normal.
2) Filter to retain only points within a depth threshold of the extremity along that direction, removing internal structures parallel to the hull.
3) Apply RANSAC plane fitting to the filtered subset.

Guiding RANSAC per direction avoids the iterative "fit and remove" strategy, which can be sensitive to extraction order and prone to fitting internal structures before hull surfaces.

D. Hopper Plate Detection

The hopper plate is a large flat surface that connects the outer side shell to the tank bottom, spanning the full longitudinal length. Its normal lies in the yz-plane with a positive z-component. We select candidate points where $z < 0$ (below the centroid), the normal has a near-zero x-component ($|n_x| < 0.1$), and the normal is not well-aligned with either the y or z axis individually ($|n_y| < 0.9$ and $|n_z| < 0.9$). These thresholds exclude points on the hull surfaces while admitting normals at the expected hopper angle of approximately 30°–50° from vertical. The candidates are further restricted to the correct transverse half-space based on the known tank side. A RANSAC plane fit on these candidates recovers the hopper plate surface.

E. Webframe Detection

The inlier points of all fitted hull surfaces and the hopper plate are removed from the cloud (using a 0.8 m distance threshold from each fitted plane) to produce an *internals cloud* containing only the internal structures. The distance threshold is chosen to be large enough to also remove smaller structure like longitudinal stiffeners that may be attached to the hull.

Webframes are vertical planar structures with normals approximately along $\pm\mathbf{e}_x$ in the tank frame. We first localize candidate webframe positions using a 1D histogram analysis, then perform targeted RANSAC at each detected peak.

1) Candidate localization: Because webframe connections to the deck are the most prominent features near the deck surface, the internals cloud is filtered to points that are (i) within 1.5 m of the detected deck plane, (ii) within a narrow transverse slice ($|y| < 0.5$ m), and (iii) have normals well-aligned with the longitudinal axis ($|\mathbf{n}_i \cdot \mathbf{e}_x| > 0.8$). Together, these filters isolate webframe-like structure near the top of the tank. The filtered points are binned along the x-axis into 100 equally spaced bins spanning the longitudinal extent of the cloud (yielding a bin width of approximately 0.25 m for a typical 25 m tank), and peaks in the resulting histogram are detected with a minimum height of 500 points.

2) Targeted plane fitting: For each detected peak at longitudinal position x_p, we gather all internals cloud points within 1.0 m of x_p whose normals are aligned with $\mathbf{e}_x$, and fit a plane to this subset with RANSAC.

3) Inlier refinement: The RANSAC inlier set may still contain stray points from internal geometry on the bottom of the tank. We apply Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [12] clustering to the inliers and retain only the largest cluster as the webframe surface.

Compared to iterative RANSAC on the full internals cloud, the histogram-based approach determines the number of webframes without requiring a known count and avoids the order-dependence problem where a large non-webframe surface could dominate early iterations. The inliers of all detected webframes are then removed from the internals cloud before the next stage.

F. Stringer Deck Detection

After webframe removal, the remaining internals cloud contains stringer decks as well as other internal structure. Since stringer decks are horizontal platforms running along the longitudinal direction at fixed heights, points with vertically aligned normals ($|\mathbf{n}_i \cdot \mathbf{e}_z| > 0.8$) are selected from the remaining internals cloud.

To reduce false detections from horizontal structure on all walls, candidates are restricted to a region near a single vertical hull surface. For each of the four vertical hull surfaces (fore and aft bulkheads, inner and outer side shells), candidates within 4 m of the wall centroid along the wall normal direction are counted. The wall with the most candidates is chosen, as it typically has the best coverage of stringer deck structure.

The selected candidates are binned along the z-axis into 100 bins and peaks are detected with a minimum height of 500 points and a minimum inter-peak distance of 3 m to avoid splitting a single stringer deck into multiple detections. For each detected peak at vertical position z_p, we fit a RANSAC plane to candidates within 0.5 m of z_p. The number and vertical positions of stringer decks follow directly from the histogram peaks.

III. Evaluation

We evaluate the proposed pipeline on accumulated point clouds from deployments of the ScoutDI 137 Gen2 Drone System [13] inside the cargo tanks of three sister crude oil tankers. We evaluate on four mid-section tanks across the three vessels: one starboard tank per vessel (V1S, V2S, V3S) and one port tank on Vessel 1 (V1P). The test clouds contain approximately 10 million points after preprocessing, and the entire pipeline runs in approximately 30 s per cloud on a laptop with an Intel Core Ultra 9 285H CPU and 32 GB of RAM. Ground truth is established by comparing detected webframe and stringer deck positions to CAD reference values. Figs. 1, 3 and 4 show qualitative results of the segmentation on one representative tank (V2S), while Table I quantitatively evaluates the detected webframe and stringer deck spacings against reference values for all tanks.

1) Hull surface and hopper plate detection: All six cardinal hull surfaces and the hopper plate were correctly detected across all four tanks. Fig. 1 shows the color-coded hull surfaces and hopper plate detections.

2) Webframe and stringer deck detection: The pipeline detected five webframes and four stringer decks in every tank. Fig. 3 shows the histogram of candidate points used for webframe and stringer deck detection as well as the detected peak locations while Fig. 4 shows the extracted webframes and stringer decks. Since the tank coordinate frame origin varies across tanks, absolute positions are not directly comparable; instead, we report inter-element spacings in Table I.

Compared to the ground truth, the detected webframe spacings match to within 0.01 m of the CAD value (4.85 m) across all four tanks, with a standard deviation of ≤ 0.005 m. The stringer deck spacings Δz_2 and Δz_3 are also well recovered, with errors ≤ 0.04 m relative to the reference values. The main outlier is Δz_1 on V1S (5.87 m vs. 5.53 m), where sparse LiDAR coverage on the topmost stringer deck leaves too few points for an accurate plane fit. The consistent detections across both starboard and port tanks (V1S vs. V1P) demonstrate generalization across tank sides.

IV. Conclusion

This paper presented a geometry-driven pipeline for semantic segmentation of cargo tank point clouds that exploited structural knowledge of mid-section cargo tanks of cargo oil tankers. Apart from the qualitative evaluation of the segmentation results, the segmented webframes and stringer decks were quantitatively evaluated against CAD reference values,

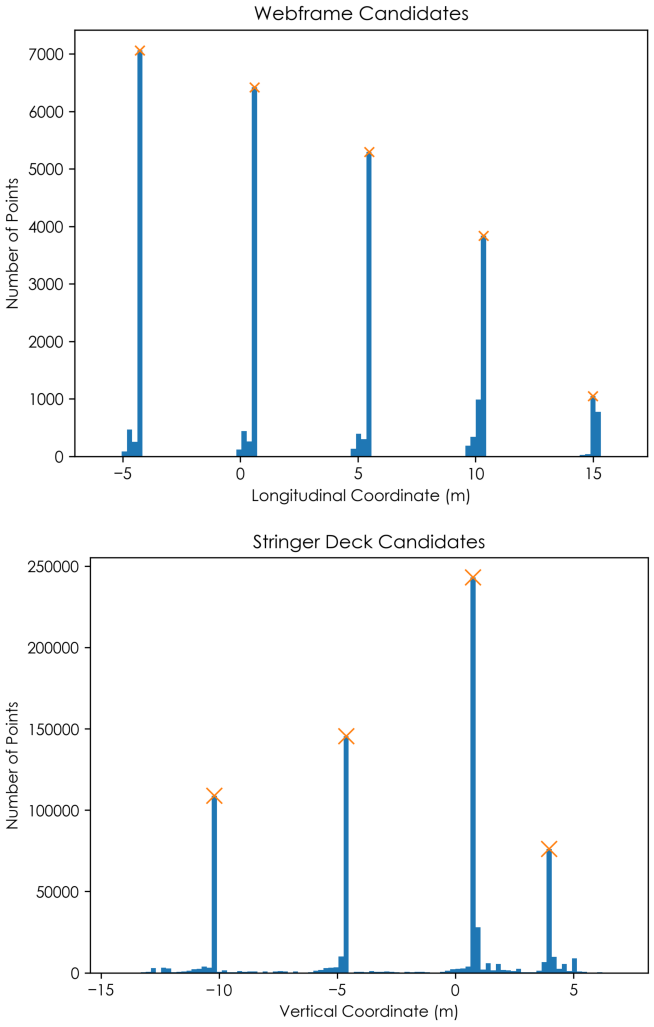

Fig. 3. Histograms of candidate points for webframe (top) and stringer deck (bottom) detection along the longitudinal and vertical axis respectively. The detected peaks are marked with × corresponding to likely locations of the respective structural elements.

TABLE I

DETECTED SPACINGS (M) ACROSS FOUR MID-SECTION CARGO TANKS ON THREE SISTER VESSELS. VnS/P DENOTES VESSEL n, STARBOARD/PORT. Δx_{1-4}: SPACINGS BETWEEN CONSECUTIVE WEBFRAMES; Δz_{1-3}: SPACINGS BETWEEN CONSECUTIVE STRINGER DECKS.

	Webframes				Stringer decks		
Vessel	Δx_1	Δx_2	Δx_3	Δx_4	Δz_1	Δz_2	Δz_3
V1S	4.85	4.85	4.85	4.85	5.87	5.50	3.19
V1P	4.86	4.85	4.85	4.86	5.50	5.49	3.23
V2S	4.86	4.86	4.85	4.85	5.51	5.50	3.18
V3S	4.85	4.85	4.84	4.85	5.47	5.51	3.17
Mean	4.86	4.85	4.85	4.85	5.59	5.50	3.19
Std. Dev.	0.005	0.003	0.003	0.003	0.163	0.006	0.025
CAD	4.85	4.85	4.85	4.85	5.53	5.53	3.2*

*As-built value verified by manual plane fitting; CAD specifies 3.42 m.

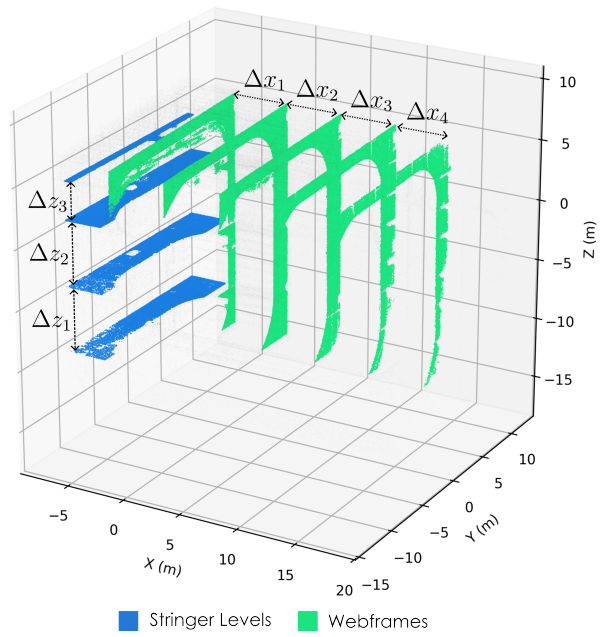

Fig. 4. Segmented webframes and stringer levels. The marked spacings between webframes (Δx_{1-4}) and stringer decks (Δz_{1-3}) correspond to the values reported in Table I.

demonstrating accurate recovery up to the noise floor of the measurements.

Aside from the restriction to mid-section tanks with the required regular structure, the main limitations are (i) sensitivity to LiDAR-inertial odometry quality, since systematic drift distorts the rectangular structure assumption and degrades hull fitting, and (ii) coverage of the internal structure by the LiDAR. Future work shall extend the pipeline to additionally segment longitudinals and to reduce the runtime for onboard execution on the inspection robot.

ACKNOWLEDGMENT

We would like to thank Maran Shuttle Tankers LLC, ScoutDI AS, and DNV AS for their support in this work.

REFERENCES

[1] M. Dharmadhikari, N. Khedekar, P. D. Petris, M. Kulkarni, M. Nissov, and K. Alexis, "Maritime vessel tank inspection using aerial robots: Experience from the field and dataset release," 2024. [Online]. Available: https://arxiv.org/abs/2404.19045

[2] R. Eckholdt Andersen, R. Yding Brogaard, and E. Boukas, "Remote inspection techniques: A review of autonomous robotic inspection for marine vessels," *IEEE Transactions on Field Robotics*, vol. 2, pp. 1–20, 2025.

[3] M. Dharmadhikari, P. De Petris, M. Kulkarni, N. Khedekar, H. Nguyen, A. E. Stene, E. Sjøvold, K. Solheim, B. Gussiaas, and K. Alexis, "Autonomous exploration and general visual inspection of ship ballast water tanks using aerial robots," in *2023 21st International Conference on Advanced Robotics (ICAR)*, 2023, pp. 409–416.

[4] R. Y. Brogaard, M. Zajaczkowski, L. Kovac, O. Ravn, and E. Boukas, "Towards uav-based absolute hierarchical localization in confined spaces," in *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2020, pp. 182–188.

[5] M. Dyduch, R. E. Andersen, and E. Boukas, "Domain adaptation method for semantic segmentation in marine vessel inspection," in *2023 IEEE International Conference on Imaging Systems and Techniques (IST)*, 2023, pp. 1–6.
[6] J. Gaetani, R. E. Andersen, and E. Boukas, "Deep stochastic image segmentation for autonomous robotic inspection," in *2023 IEEE International Conference on Imaging Systems and Techniques (IST)*, 2023, pp. 1–6.
[7] R. E. Andersen and E. Boukas, "Probabilistic inspection using deep probabilistic segmentation," in *2022 IEEE International Conference on Imaging Systems and Techniques (IST)*, 2022, pp. 1–6.
[8] M. Waszak, A. Cardaillac, B. Elvesæter, F. Rødølen, and M. Ludvigsen, "Semantic segmentation in underwater ship inspections: Benchmark and data set," *IEEE Journal of Oceanic Engineering*, vol. 48, no. 2, pp. 462–473, 2023.
[9] Z. Li and J. Shan, "RANSAC-based multi primitive building reconstruction from 3D point clouds," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 185, pp. 247–260, 2022.
[10] S. Tang, J. Huang, H. Du, B. Zhou, Z. Zhao, W. Wang, and R. Guo, "Back to geometry: Efficient indoor space segmentation from point clouds by 2D–3D geometry constrains," *International Journal of Applied Earth Observations and Geoinformation*, vol. 135, p. 104265, 2024.
[11] N. Khedekar and K. Alexis, "PG-LIO: Photometric-Geometric fusion for Robust LiDAR-Inertial Odometry," Jun. 2025, arXiv:2506.18583 [cs]. [Online]. Available: http://arxiv.org/abs/2506.18583
[12] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, ser. KDD'96. AAAI Press, 1996, p. 226–231.
[13] "Scout 137 Drone System — scoutdi.com," https://www.scoutdi.com/scout-137-drone-system/, [Accessed 23-03-2026].