

Constructing Visibility Maps of Optimal Positions for Robotic Inspection in Ultra-High Voltage Centers

Panagiotis Mermigkas*, George P. Moustris, Costas S. Tzafestas and Petros Maragos

Abstract—Visibility maps are crucial for autonomous robotic applications such as exploration, path planning, obstacle avoidance, and multi-robot coordination. In the context of electrical transmission infrastructure, automated robotic inspection enhances proactive maintenance, enabling early detection of wear, damage, or faults, thereby improving safety, extending component lifespan, and optimizing maintenance schedules.

In this work, we propose an algorithm to compute optimal visibility locations, enabling a mobile robot to acquire RGB and thermal images for fault detection. Using LiDAR scans, we construct a global 3D map composed of ground structures (represented as a Grid Map) and overground structures (modeled with an OctoMap for efficient ray-casting). We apply clustering techniques to identify 3D bounding boxes for electrical components and define suitable source and target points for visibility assessment. By employing a weighted visibility scoring approach, we determine the ground positions that offer the best visibility of each component while ensuring minimal occlusions and adherence to viewing constraints.

The proposed method enables a robot to autonomously navigate to these optimal viewpoints, improving inspection efficiency. By integrating visibility regions across multiple components, the inspection process is further optimized, reducing overall inspection time. Our algorithm has been successfully deployed and tested at an Ultra-High Voltage Center (UHVC) in Greece, demonstrating its effectiveness in real-world conditions.

I. INTRODUCTION

Automated robotic inspection of critical infrastructure presents a valuable solution for maintaining continuous operation, improving uptime, and extending the lifespan of these systems. Performed in demanding or hazardous environments, it reduces risks for human inspectors while enhancing the efficiency of the inspection process. A key application of this technology is in electrical transmission infrastructure, where precise monitoring of essential components is crucial. By automating inspections and utilizing both visible-spectrum (RGB) and infrared imaging, power operators can thoroughly assess component conditions, identifying wear, damage, or faults early. This proactive approach helps prevent failures, improves safety and reliability, optimizes maintenance schedules, and reduces costs.

This work originates from the ENORASI project (enorasi-insight.com), which focuses on the automated inspection of Ultra-High Voltage Centers using autonomous ground vehicles. In a typical scenario, a mobile robot is assigned a set of specific electrical components requiring inspection. These components are selected based on an inspection policy that considers factors such as optical or thermal wear, the



Fig. 1: Employment and testing of our proposed algorithm at a UHVC in Greece (right). Our mobile robot navigates to one of the optimally selected locations that exhibits a high visibility score for the electrical component under inspection (top left, Voltage Transformer shown in the green 3D bounding box), where our Viewpro Z10TIR RGB and thermal camera captures images for inspection and fault evaluation (bottom left).

time elapsed since the last inspection, and the component's criticality. To conduct the inspection, a visitation plan is computed using a multi-criterion planning algorithm. This planner considers a set of "optimal" visibility points for each component and produces a path that enables the robot to navigate through them, in order to capture RGB and thermal images (see Fig. 1).

Our overall system enables full automation of the inspection rounds, considering ground conditions such as rough terrain, holes, and obstacles, as well as optimizing the robot's positioning for image acquisition. By integrating visibility regions from multiple inspection elements, the process is further optimized, allowing the robot to navigate to vantage points where multiple targets are visible. This reduces inspection time and improves overall system efficiency. The robot autonomously calculates the optimal visitation plan and positioning while dynamically adapting to environmental changes.

In this paper, we present an algorithm for constructing visibility maps for this inspection process. Using a precomputed 3D map of the UHVC, we generate a voxel-based

All authors are with the School of Electrical & Computer Engineering, National Technical University of Athens, Greece, Corresp. author email: *p.mermigkas@gmail.com

representation of the electrical components along with a suitable elevation model of the terrain. For each component, we extract its bounding box and, through a ray-casting process, identify ground positions that offer “optimal” visibility for each face. These positions allow clear observation of points-of-interest on each component, without occlusions and adhering to various constraints, such as the viewing angle. The highest scoring positions are then clustered together, forming a neighborhood of candidate visibility points used in the planning algorithm. In the following, we present a detailed analysis of the various steps and algorithms of the entire pipeline. For the benefit of the community, we provide the open source of our implementation at https://github.com/pmermigkas/octomap_visibility_uhvc.

II. RELATED WORK

Determining optimal camera viewpoints that maximize visual information or minimize uncertainty is a fundamental principle in the domains of object reconstruction, scene reconstruction, object recognition, and pose estimation [1]. To deal with uncertainty in sensor positioning, [2] re-evaluated the initially selected candidate views according to their neighboring views, ensuring that the best view is located within a region of high-quality viewpoints. In our work, we compute visibility scores for each cell of a 2D Grid Map, using a morphological erosion step to mitigate the negative effects of uncertainties in sensing, positioning and registration.

In [3], a set of robot-reachable viewpoints is evaluated through a surroundings check to eliminate those obstructed by obstacles or occlusions that hinder visibility of the candidate object. Subsequently, a model analysis is conducted to assess the recognition probability at each feasible viewpoint. The next-best view is selected as the viewpoint that maximizes both visibility and recognition criteria. Similarly, in [4], the authors determine the next-best view by employing a ray-casting procedure on octrees to exclude occluded object features. This process is iteratively repeated until robust object recognition and pose estimation convergence are achieved. In [5], the next-best view is determined by performing ray-casting to identify visible voxels, enhanced by a quality criterion based on the computation of surface normals within a small region around each image pixel.

In contrast to traditional ray-casting algorithms, [6] introduce a lightweight dynamic-programming framework for solving 2D grid visibility problems. Their approach leverages a linear hyperbolic partial differential equation (PDE) to propagate visibility in all directions across a 2D grid, ensuring convergence to the true visibility polygon as grid resolution increases. Another work offering an approximate solution for visibility map computation is [7], which proposes a computationally efficient alternative to the brute-force approach for robots of any shape. In contrast, our pipeline employs exact ray-casting methods, prioritizing high robustness to sensor noise as its most critical feature.

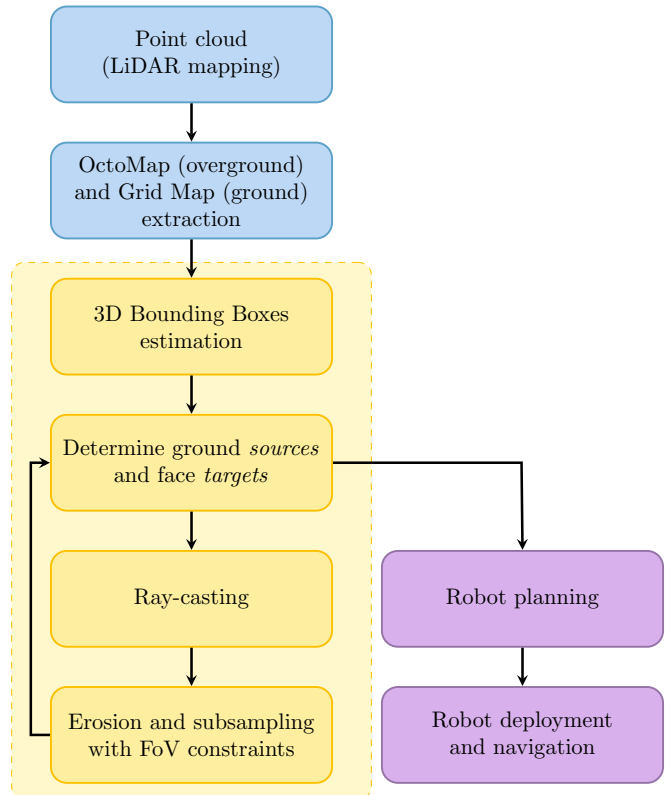


Fig. 2: Our proposed system overview includes a pre-processing step (blue), the main ray-casting loop for computation of visibility scores (yellow), and the real world deployment (purple) step.

III. METHODS

In the following sections, we provide a detailed analysis of the pre-processing and the main loop components of our system pipeline, which are highlighted in blue and yellow, respectively, in Fig. 2.

A. Initial Map Construction

In order to construct the visibility regions, we first need to create a 3D map of the UHVC. To this end, we used a mobile robot (Summit XL, Robotnik Automation S.L.) equipped with a Light Detection and Ranging (LiDAR) sensor running a Simultaneous Localization and Mapping algorithm based on FAST-LIO2 [8]. Using teleoperation, we navigated a significant portion of the area, capturing key regions of interest. Following data collection and initial processing, we constructed a point cloud representation of the environment.

However, the resulting map contained significant noise due to sensor and algorithmic limitations, leading to excessive surface thickness that interfered with key autonomous inspection methods such as ray-casting. To address this issue, we applied a two-step filtering process; first, we used a uniform sampling filter to assign voxels to the 3D space, retaining only the closest voxel for each point. This reduced noise while preserving essential structural information. Next, we applied the Moving Least Squares (MLS) algorithm to refine




Fig. 3: View of the UHVC elevation map, overlaid with the point cloud of the overhead structures. The ground is colored according to its respective cell height.

the point cloud by fitting local surfaces, which diminished noise and improved normal estimation. While this method slightly smoothed sharp edges, it significantly enhanced the overall quality of the point cloud. More details regarding the map creation process can be found in [9].

B. Ground Model

Creating a model of the ground is crucial, as the *visibility map* itself uses this data structure to denote visible regions. As a first step, we segmented the ground points in the point cloud from the overhead structures, using a *cloth simulation filter* (CSF) with appropriate parametrization [10]. The extracted points were then used to create an *elevation map* using the Grid Map library [11]. In this format, the ground is divided into square cells of the same size, where each cell encodes various information about the points within it. For example, each cell can encode elevation, the normal vector, roughness, traversability etc. For the elevation, we used the corresponding filter, computing a single elevation value for each cell from the coordinates of its included points and the respective normal filter to create the Grid Map layer representing the ground normals to each cell (Fig. 3). By removing the ground points from the overall point cloud, we can extract the points of the *overhead structures*, which include the electrical components for inspection.

C. Voxel-Based Representation of Components

To assess the visibility of each electrical component (EC), we need to represent it as a volumetric object in 3D space. Thus, the description of overground structures using point cloud data must be processed to create a structured representation that accurately defines their shape and spatial occupancy. To this end, we have opted for using the OctoMap framework, based on the *octree* data structure [12]. In this case, the three-dimensional space is discretized into volumetric elements (voxels), forming a hierarchical structured representation that classifies regions as occupied, free, or unknown.




Fig. 4: View of the OctoMap of the overground structures at the UHVC.

OctoMap offers several advantages for ray-casting, making it highly effective for visibility analysis and sensor simulation in robotics. Its octree-based spatial representation enables efficient storage and retrieval of occupancy information at multiple resolutions, allowing for fast ray traversal by skipping large empty or occupied regions. This significantly reduces computational overhead compared to checking every individual point in a dense point cloud. Additionally, OctoMap supports multi-resolution adaptability, enabling quick ray-casting in large environments using coarse representations while maintaining finer details for precise visibility computations. It also allows efficient updates in dynamic environments, ensuring accurate ray-casting even as objects move or change. Furthermore, it is memory-efficient as its structure is compact, making it suitable for large-scale environments without excessive computational costs. A view of the OctoMap of the UHVC can be seen in Fig. 4.

D. 3D Bounding Boxes extraction

The ECs of interest are Voltage Transformers (VTs), Current Transformers (CTs), Bushings (BSs), Surge Arresters (SAs), Isolators (ISs) and Circuit Breakers (CBs), which are depicted throughout Figs. 3 and 4. To achieve this annotation, we leverage the fact that all ECs are positioned above the height of 3.5 meters from the ground level, in order to discard the supporting base metal parts and therefore aiding the clustering significantly. More importantly, the remaining part of the point cloud is fully symmetrical in the z -axis due to the cylindrical shape of all ECs, which reduces the clustering process to a two-dimensional space. Next, we run k -means clustering, resulting in a set of *cluster centers* and *cluster members*. The final 3D bounding boxes are defined by the enclosing cuboid of the 3D points of each cluster (see Fig. 6 for a visualization of the bounding boxes' locations and sizes). It should be noted that the UHVC global map comprises ECs which are organized in a grid-like structure aligned to the global xy axes. Consequently, the bounding boxes are also inherently axis-aligned with the global coordinate system, ensuring computational convenience and consistency.




Fig. 5: 3D visualization of visibility regions underneath an electrical component represented as a 3D bounding box (cuboid with a blue face) before (bottom, red layer) and after (top, green layer) the application of the erosion step. The gray cube represents a blocking element that restricts vision towards six out of nine ray-casting targets, leading to lower visibility scores compared to the leftmost source.

E. Determination of view count

An important constraint in our use case is that the RGB and thermal images must be captured from various angles to detect potential thermal or optical faults across the entire EC surface. To satisfy this requirement, it was determined that each EC inspection should include four images, each taken from a side of the grid-like UHVC structure. This approach is naturally aligned with the narrow corridors of UHVCs and ensures comprehensive EC coverage with the minimum number of images required.

F. Determination of ray-casting sources & targets

For each side of the aforementioned cuboids, feasible visibility regions inherently exclude the area behind that face, as well as regions located far from the cuboid, as the captured ECs would appear too small in the image. Consequently, we have opted to consider possible visibility regions in a rectangular area with dimensions $W \times H$ on the ground directly in front of the EC. This region is discretized into cells using the same spatial resolution as the OctoMap voxel grid, denoted by D_V , which was set to 4cm in our implementation; however, this value can be adjusted to any arbitrary spatial resolution as needed. Fig. 5 illustrates this region as the bottom (red) 2D grid. The *sources* for ray-casting are located at positions

$$\mathbf{x}_s = (x_s, y_s, z_s)^\top + H\mathbf{n} \quad (1)$$

where x_s and y_s correspond to the center coordinates of each cell, z_s represents the respective elevation value, H is the robot camera height and \mathbf{n} is the ground normal.

To evaluate the visibility of a ground cell, we perform ray-casting towards one of the cuboid faces. More specifically,

we overlay a $(n-1) \times (n-1)$ grid on the cuboid face, defining a set of $N_f = n^2$ ray-casting *targets* located at the corners and intersections of this grid, as illustrated with the blue rectangle in Fig 5 in the case of $n = 3$. By performing ray-casting from all ground cells towards all N_f points, for a total of WHN_f/D_V^2 ray-casting operations, we assess the visibility of the entire face. Based on the number of visible target points, we assign a *visibility score* to each ground cell.

G. Weighted visibility scoring

Given that the top part of the cuboid face typically contains the majority of points of interest for our inspection tasks, we employ a weighted approach that assigns higher weights to the upper section of each face. This is due to the fact that electrical connections, which are the most likely points of failure, are predominantly located in the mid-to-upper section of each EC. Finally, we introduce the visibility scores, v_{ij} , for ray-casting from a source point at position $\{i, j\}$ within the 2D ground grid as:

$$v_{ij} = \sum_{k=1}^{N_f} \mathbf{1}_{ij}(k) w_k \quad (2)$$

where w_k represents the individual weights assigned to each face target, and $\mathbf{1}_{ij}(k)$ denotes the indicator function of the subset of visible face targets at position $\{i, j\}$.

H. Erosion step

During the creation and post-processing of point clouds and OctoMaps, structural gaps or noise are frequently observed, leading to either overestimated or underestimated visibility evaluations, respectively. Consequently, these derived visibility regions often exhibit isolated “valleys” and “hills” that contradict the intuitive human perception of 3D visibility, where smooth and continuous visibility transitions are expected.

To address these issues, we implement a post-processing step consisting of a minimization operation followed by a subsampling step. First, we apply the “minimum operator” within a square area centered on each ground cell, which can be thought of as the analogous to performing a 2D morphological erosion on an image. Next, we perform subsampling at a lower spatial resolution, denoted by D_S . The resulting 2D visibility grid \tilde{v} is visualized as the top (green) layer in Fig. 5. This procedure of obtaining the processed visibility scores $\tilde{v}_{i'j'}$ at position $\{i', j'\}$ and resolution D_S can be formally described by the following equation:

$$\tilde{v}_{i'j'} = \min_{\{i,j\} \in \text{Sq}} v_{ij} \quad (3)$$

where Sq represents a square area around the desired position $\{i', j'\}$ and contains all points with the property

$$\left\| \left(i' + \frac{1}{2} \right) D_S - \left(i + \frac{1}{2} \right) D_V \right\|_\infty \leq D_S \quad (4)$$

It should be noted that we have assumed that the centers of the xy coordinate systems are located at the top left corner of each grid as illustrated in Fig. 5.




Fig. 6: Procedure for the computation of visibility scores for the case of the back face of a *Bushing (Phase B)* electrical component (illustrated as a light green 3D bounding box): visibility scores at a spatial resolution of $D_V = 4cm$ (left), visibility scores after applying the FoV constraints (middle), and final visibility scores after the erosion and subsampling step at a spatial resolution of $D_S = 15cm$ (right). Brighter (darker) areas correspond to locations that exhibit greater (lower) visibility towards the back face of the EC. Small black squares correspond to ground anomalies or holes, which are not traversable for the robot vehicle.

I. Constraints on the field of view (FoV)

After applying the previous steps, we obtain visibility scores for each cell in front of each face, resulting in a total of $\left\lceil \frac{W}{D_S} \right\rceil \cdot \left\lceil \frac{H}{D_S} \right\rceil$ cell values. These scores quantify the presence of occlusions between the source and the N_f face targets without accounting for the constraints imposed by the thermal camera's FoV.

More specifically, due to the projective camera geometry there is a critical point beyond which, as the sensor moves horizontally towards the EC, the FoV frustum no longer contains the entire EC, which is a mandatory requirement for UHVC inspection. Consequently, it is necessary to identify regions where the thermal camera fails to capture all N_f points and exclude them from the visibility analysis to ensure the desired inspection results.

Another constraint arises from the limitations of the mounted robotic platform, since the sensor is mounted on a platform that physically obstructs the view directly above and behind it. As a result, there is a maximum attainable pitch angle for the camera. In the yaw direction there is also a range of angles which block the view, but since the robot can freely rotate around its vertical axis, this constraint can be effectively disregarded in practice. The effect of this constraint is demonstrated by the semi-circular pattern formed in the middle image of Fig. 6.

Finally, we have opted to exclude regions at viewing angles close to 90° with respect to the orientation of the face normal, because the EC under inspection projects to the camera as an extremely narrow object from those viewpoints.

IV. EXPERIMENTAL EVALUATION

To collect the data for map creation and evaluate the efficacy of our proposed algorithm, we deploy our mobile robot, equipped with a Robosense RS-LiDAR-16 and a

Vectornav VN-100 IMU, at a UHVC in Greece, which contains hundreds of ECs. Following data collection, we obtain and post-process a static denoised map of the UHVC according to the methods detailed in [9]. Next, according to the methodology analyzed in the previous section, we developed all components of our system pipeline in a novel C++ ROS node, in order to compute visibility scores for the most critical ECs of a UHVC. For the scope of this work, we focused on extracting visibility scores for a test area that contains two gates with all six EC classes of interest (VTs, CTs, BSs, SAs, ISs and CBs). We compute and provide the full set of visibility score results in our code repository for a total of 42 ECs (gate A has eight rows and gate B six rows, each multiplied by the three phases), along with the OctoMap representation to aid in the visualization and interpretation of results. Four visibility regions correspond to each EC, one for each face of the cuboid bounding box that encloses it (see Fig. 7 for a sample visualization of the resulting visibility regions for a *Bushing (Phase B)* EC).

The most critical intermediate results of our procedure are illustrated in Fig. 6. More specifically, in the leftmost image we show visibility scores for the case of the back face of a *Bushing (Phase B)* EC. The 2D grid spans a ground area of 15 meters \times 7 meters at a spatial resolution of $D_V = 4cm$. The weights w_k used are equal to 1, 3, and 9 for the bottom, middle and top rows of the ray-casting targets, respectively, leading to a maximum attainable score of 39 (brighter regions) and a minimum score of 0 (darker regions). In the middle image, we have applied the FoV constraints for the thermal camera in the vertical and horizontal directions, and excluded locations at angles greater than 80° relative to the face normal direction. Finally, in the rightmost image, we have applied a minimization and subsampling step to provide robustness and resistance to noisy structures.




Fig. 7: Results for the final visibility regions for all four faces (back, left, front, right in order of appearance from left to right) of a *Bushing (Phase B)* electrical component (illustrated as a light green 3D bounding box). The top row corresponds to the case where $w_k = 1$ for all ray-casting targets, while the bottom row corresponds to a weighting scheme of $w_k = 1, 3$, and 9 for the bottom, middle and top rows of the ray-casting targets, respectively.

V. CONCLUSION

In this work, we presented a novel algorithm for constructing visibility maps to identify optimal positions for autonomous robotic inspection in UHVCs. Our approach leverages LiDAR-based mapping to generate a voxel-based OctoMap representation of overground ECs and a Grid Map for ground elevation modeling. By applying clustering techniques, ray-casting, a weighted visibility scoring system, and a morphological erosion operator, we determine error-tolerant ground positions with minimal occlusions, ensuring comprehensive RGB and thermal imaging for fault detection.

The proposed method enables efficient and automated inspection by allowing a mobile robot to navigate through precomputed optimal viewpoints, reducing inspection time while maximizing coverage. By integrating visibility regions across multiple components, our approach further optimizes the inspection process, enhancing reliability and minimizing the need for manual intervention. Additionally, our methodology accounts for real-world constraints, such as sensor limitations and environmental obstructions, ensuring robust performance in practical deployment scenarios.

We validated our approach through real-world experiments at a UHVC in Greece, demonstrating its effectiveness in identifying optimal viewpoints for fault detection. Future work will focus on improving real-time adaptability in dynamic environments, incorporating machine learning-based optimization for viewpoint selection, stacking multiple bounding boxes vertically to accommodate co-located elements at different heights, and expanding the methodology to other industrial inspection domains.

VI. ACKNOWLEDGEMENT

This research has been co-financed by the European Union NextGenerationEU under the call RESEARCH – CREATE – INNOVATE 16971 Recovery and Resilience Facility (project: ENORASI, code: TAEDK-06172).

REFERENCES

- [1] R. Zeng, Y.-H. Wen, W. Zhao, and Y.-J. Liu, “View planning in robot active vision: A survey of systems, algorithms, and applications,” *Computational Visual Media*, vol. 6, no. 3, pp. 225–245, Sep. 2020.
- [2] J. I. Vazquez-Gomez, L. E. Sucar, R. Murrieta-Cid, and E. Lopez-Damian, “Volumetric Next-best-view Planning for 3D Object Reconstruction with Positioning Error,” *Intl. Journal of Advanced Robotic Systems*, vol. 11, no. 10, Oct. 2014.
- [3] C. McGreavy, L. Kunze, and N. Hawes, “Next Best View Planning for Object Recognition in Mobile Robotics,” in *Proc. of Workshop of the UK Planning and Scheduling Special Interest Group*, Huddersfield, England, Dec. 2016.
- [4] K. Wu, R. Ranasinghe, and G. Dissanayake, “Active recognition and pose estimation of household objects in clutter,” in *Proc. of IEEE Intl. Conf. on Robotics and Automation*, Seattle (WA), USA, May 2015.
- [5] N. A. Massios and R. B. Fisher, “A Best Next View selection algorithm incorporating a quality criterion,” in *Proc. of the British Machine Vision Conf.*, Southampton, UK, Sep. 1998.
- [6] I. Ibrahim, J. Gillis, W. Decré, and J. Swevers, “An Efficient Solution to the 2D Visibility Problem in Cartesian Grid Maps and its Application in Heuristic Path Planning,” in *Proc. of IEEE Intl. Conf. on Robotics and Automation*, Yokohama, Japan, May 2024.
- [7] T. Pereira, M. Veloso, and A. Moreira, “Visibility maps for any-shape robots,” in *Proc. of IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, Daejeon, Korea, Oct. 2016.
- [8] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, “FAST-LIO2: Fast Direct LiDAR-Inertial Odometry,” *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, Aug. 2022.
- [9] I. Alamanos, G. P. Moustris, and C. S. Tzafestas, “Localization and Offline Mapping of High-Voltage Substations in Rough Terrain Using a Ground Vehicle,” in *Proc. of 32nd Mediterranean Conf. on Control and Automation*, Chania, Greece, Jun. 2024.
- [10] W. Zhang, J. Qi, P. Wan, H. Wang, D. Xie, X. Wang, and G. Yan, “An Easy-to-Use Airborne LiDAR Data Filtering Method Based on Cloth Simulation,” *Remote Sensing*, vol. 8, no. 6, Jun. 2016.
- [11] P. Fankhauser and M. Hutter, “A Universal Grid Map Library: Implementation and Use Case for Rough Terrain Navigation,” *Robot Operating System (ROS) – The Complete Reference (Volume 1)*, pp. 99–120, Feb. 2016.
- [12] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, “OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees,” *Autonomous Robots*, vol. 34, no. 3, pp. 189–206, Apr. 2013.