# An autonomous inspection robot for extra-high voltage substations *

George Moustris[1], Thaleia Flessa[2], Panagiotis Mermigkas[1], Ioannis Alamanos[1], Nikolaos Kardaris[1], Andreas Vatistas[1], Costas Tzafestas[1], Petros Maragos[1], George Bonias[2], Angelos Skembris[2], Aristeidis Geladaris[3,4], Athanasios Mastrogeorgiou [3,5], Jun Rong[6], Panagiotis Bagiokos[6], Despoina Makrygiorgou[6], Serafeim Panidis[6], Dimitris Melissaris[6]

[1]*School of Electrical and Computer Engineering*, *National Technical University of Athens*, Greece
[4]*Mechanical Engineering Dept*, *Hellenic Mediterranean University*, Heraklion, Greece
[5]*School of Mechanical Engineering*, *National Technical University of Athens*, Greece
[6]*Department of Research Technology and Development*, *Independent Power Transmission Operator*, Greece
[2]*Castalia Technologies SMPC*, Athens, Greece
[3]*Tech Hive Labs*, Athens, Greece

***Abstract*—Extra High Voltage Centers are critical to power transmission, yet their maintenance still relies on manual, time-based inspections with handheld thermal cameras. This approach is labor-intensive, costly, and prone to missing early faults. In this paper we present ENORASI, an autonomous robotic platform for condition-based inspection of EHVCs. ENORASI integrates 3D LiDAR-based mapping, visibility scoring, and traversability analysis to enable safe and efficient navigation in complex substation terrain. A decision support system enables smart scheduling and report generation, while an optimal path planner solves a Generalized Traveling Salesman Problem for efficient routing. Component faults are detected by combining RGB-based object recognition with thermal image analysis. Initial deployments demonstrate that ENORASI can increase inspection frequency, reduce operational costs, and support the transition from time-based to condition-based maintenance. These results validate ENORASI as a scalable and effective alternative to manual inspection for critical power infrastructure.**

***Index Terms*—autonomous robot, inspection, extra-high voltage substation, planning, thermal imaging**

## I. Introduction

High and Extra High Voltage Centers (EHVC) are essential nodes in the power system, serving as key hubs for the transmission of electricity across the country. EHVCs are crucial in managing and interconnecting the electricity generation sources with the points of consumption, whether residential or industrial. The maintenance of EHVCs is under the responsibility of the Independent Power Transmission Operator (IPTO). Their operation is crucial to ensure the safe, uninterrupted and reliable power supply and to facilitate grid resilience during faults or emergencies. Failures due to material defects or transient events can lead to major disruption and outages as seen, for example, during a 2018 incident where a Circuit Breaker (CB) in Pallini's EHVC presented an unexpected failure, which caused power outages in large parts of Athens, which is home to nearly half of the population of Greece, and lead to major societal impacts such as transportation disruption, and affected public buildings [1].

Currently, maintenance is time-based, following manufacturer schedules. This approach can be risky, as components may fail before their scheduled replacement, as was the case for the new circuit breaker at the Pallini EHVC that failed before its 2020 service date causing the 2018 outage. Time-based maintenance may also lead to unnecessary replacements, raising costs, especially when operating conditions deviate from the manufacturer's assumptions—something exacerbated by climate change. A more effective approach is condition-based maintenance, which relies on real-time monitoring to assess their actual health. Electrical damage isn't necessarily visually detectable and thermal imaging is the most effective diagnostic method. Specially trained IPTO personnel use handheld thermal cameras for periodic checks, but this is insufficient for systematic monitoring and there is a lack of available technology to enable frequent, automated measurements.

Static thermal and optical cameras are not ideal due to high installation costs and the need for repositioning when equipment changes. Transitioning to condition-based maintenance requires advanced, automated technologies for data capture, and analysis. Automated inspection methods can fill in this gap by providing a more flexible, robust and cost-effective solution. A prominent example is robotic inspection, where a robot is used to perform scheduled inspection in EHVCs, at a significantly higher frequency than manual inspection. These robots can navigate complex environments, capture high-resolution thermal and visual data, and upload results for real-time analysis, enabling early fault detection, reducing downtime, and improving system reliability without the need for human presence. The final decision on necessary actions

*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).

979-8-3315-2503-3/25/$31.00 ©2025 IEEE

Proc. of 2025 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) Conference, Valletta, Malta, October 20-23, 2025.

Fig. 1. View of the inspection robot inside an EHVC

(e.g., replacement) will still rest with the human operators.

This work stems from the ENORASI project, which targets the automated inspection of EHVCs using autonomous ground vehicles (Fig.1). In a typical deployment, a mobile robotic vehicle is tasked with inspecting a predefined set of electrical components. These components are selected according to an inspection policy that takes into account parameters such as thermal or optical degradation, last inspection date, and component criticality. The robot autonomously plans the optimal inspection route, navigates through the EHVC, acquires optical and thermal images of the components, which are then analyzed to identify defects. In the following sections, we present the overall system architecture and the main components comprising our solution.

## II. System Overview

The ENORASI platform consists of the robotic vehicle, a server (to store images, data and run demanding algorithms such as the path planning and defect detection) and a user interface control center, where the operator has an overview of the inspection process. The operator can access information on the electrical components and inspection history and inspections are planned automatically. An inspection plan can also be initiated on demand by the user. The entire pipeline is divided into four distinct components and processes, each facilitating a specific task: *Map Creation*; *Decision Support System*; *Inspection Planning*; *Inspection Execution*. Each one is analyzed in the following sections.

## III. Map Creation

### A. Mapping of the Inspection Site

To generate the initial map, we utilize the FAST-LIO2 [2] framework for estimating poses and generating 3D point clouds. The raw LiDAR measurements are first motion-compensated using scan undistortion. This is achieved by back-propagating each point based on inertial data from the IMU. Pose estimation is then performed for each LiDAR scan. This involves formulating a non-linear least squares problem, to minimize residuals arising from aligning the current scan to the existing map using an Extended Kalman Filter.

However, due to sensor noise and algorithmic limitations, the resulting map, comprising accumulated scans registered to the odometry frame, exhibits artifacts, including surface thickening and noise. This adversely affects operations such as raycasting between inspection targets and the ground. To mitigate this, we apply a two-stage filtering pipeline;

We first apply a voxel-based uniform sampling filter, which discretizes the 3D space and retains only one point per voxel. This not only reduces noise but enforces spatial regularity, resulting in a sparser and more structured point cloud. Subsequently, we apply a Moving Least Squares (MLS) algorithm to smooth and de-noise the cloud. MLS fits local surfaces by minimizing a weighted least squares cost:

$$\sum_{i=1}^{N}(\langle n_i p_i\rangle - D)^2\theta(||p_i - q||), \tag{1}$$

where the local plane is defined as: $\kappa = x|\langle n_i p_i\rangle - D = 0, x \in \mathbb{R}^3,\ n \in \mathbb{R}^3,\ ||n|| = 1$, $p_i$ is a 3D point, $n_i$ its corresponding normal, $D$ the plane model, q the projection onto the plane, and $\theta$ a smooth monotonous decreasing function. This smoothing process enhances normal estimation and significantly suppresses noise. Although MLS tends to slightly smooth sharp features, the overall denoising outcome is highly beneficial for downstream tasks.

### B. Visibility Maps

To perform inspection tasks, the robot must be aware of optimal map positions, which provide the least occluded viewing angles for capturing thermal images. These positions constitute a *visibility map*, assigning a score to every position in a 2D grid on the ground, enabling view planning strategies that prioritize safety, coverage, and efficiency.

Given the preprocessed 3D point cloud of the inspection site, we extract the terrain by applying a Cloth Simulation Filter [3] and convert overground points into an Octomap with a voxel resolution of 5 cm. This hierarchical voxel-based map provides an efficient occupancy representation with occupied voxels representing occlusions and supports fast ray-casting which significantly accelerates visibility computation.

Based on [4], to assess visibility for each electrical component (EC), we first extract axis-aligned 3D bounding boxes via HDBSCAN 2D clustering, exploiting the symmetry of near-cylindrical ECs and discarding the lower metal parts of the structures. These bounding boxes serve as proxies for real geometry and sampled points on their four vertical faces are utilized as *ray-casting targets* for visibility scoring.

To define *ray-casting sources*, we consider a 2D ground grid of resolution 4 cm over the robot-navigable area. At each cell, we cast rays from a virtual observer at a fixed height of 75 cm above ground toward the aforementioned ray-casting targets. Finally, we introduce *visibility scores* as the normalized amount of rays reaching the target volume, resulting in scalar values between 0 (fully occluded) and 1

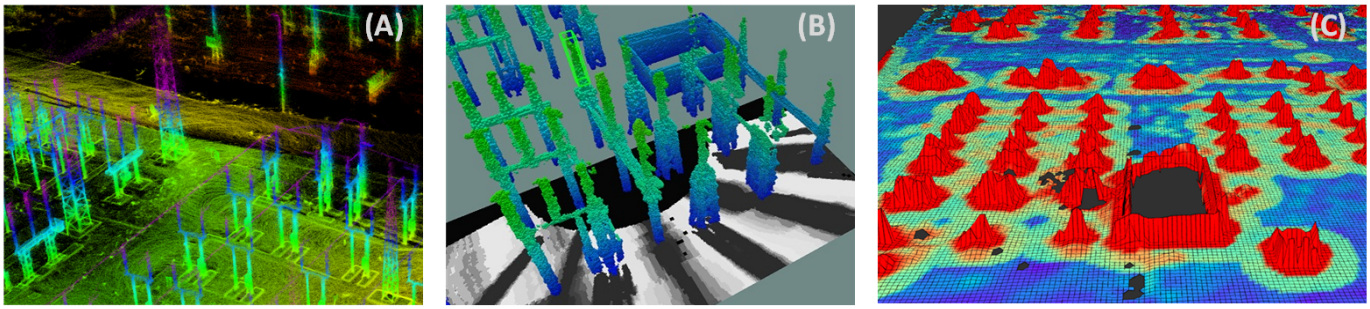


Fig. 2. A) 3D Point cloud of inspection sit. B) Visibility maps of various components. C) Traversability map of the inspection site.

(fully visible). This process produces one visibility map per EC, as shown in examples of Fig. 2(B) for various ECs, revealing spatial regions that maximize inspection quality.

### C. Traversability Map

Determining which parts of the rough terrain in an EHVC can be traversed is crucial. Based on the previously described maps, we use a CSF to isolate ground from non-ground segments. Since the vehicle's height is slightly less than 1.5 meters, we keep only the ground and points up to 1.5 meters above it. This strategy excludes higher areas where collisions cannot occur and focuses the traversability analysis on the vehicle's operating environment.

Next, we employ a 2D grid to encode the segmented terrain. The resulting point cloud is represented as a grid of cells in which each cell contains an *elevation* value. To determine which cells are traversable, we apply three additional filters on their respective layers, viz. a *Surface Normals Filter*, computing a normal vector for each cell, a *Slope Filter*, calculating the slope of each cell and a *Roughness Filter*, quantifying the unevenness of each cell by examining its surrounding normals. By normalizing both the *Slope* and *Roughness* layers and applying thresholds that reflect the vehicle's terrain-handling abilities, each cell is assigned a *traversability* cost (Fig. 2(C)). This is subsequently incorporated into ROS for planning and navigation purposes.

## IV. Decision Support System

### A. Overview

The decision support system's (DSS) primary function is to enable automated inspections of EHVC equipment and provide access to historical inspection data. The aim is to develop a user-friendly system for scheduling inspections, monitoring equipment status, accessing related information, generating automated reports, and also integrating with other components like path planning and defect detection. The system consists of software running on the robotic vehicle (local path planning, navigation, image acquisition), server-based software (path planning, defect detection) and a database, also stored on the server. All data and images are stored in this database and an accompanying image repository. The user interface (UI) is hosted on the server, allowing users to securely access it via personal computers to view past plans, schedule new inspections, and generate reports. Users can also access data from previous inspections, including optical and thermal images and the condition of components determined by the defect detection algorithm. Each component's data includes identifying information (type, name, gate, electrical phase, installation date), reliability indices based on image analysis assessments of thermal and optical degradation, inspection history, and links to the image repository. This comprehensive dataset allows for informed decision-making regarding inspection scheduling and maintenance planning.

### B. Inspection Scheduling

The DSS includes a task scheduling framework [5] that generates EHVC inspection plans and outputs a list of components for inspection, which is passed to the path planning system to compute an optimal navigation path. During inspections, thermal and optical images are analyzed to assign degradation indices to each component. These indices, along with the component's age and criticality, are used to calculate the Total Reliability Index (TRI) [5], which reflects the probability of failure and guides the frequency of monitoring.

There are three types of inspection plans [5]. The automated monthly plan runs without user input and includes all components, representing a time-based maintenance strategy [6]. The weekly plan is condition-based [7] and targets components with a TRI below a predefined threshold. The user-defined plan allows users to select components based on criteria such as type, age, TRI, thermal or optical damage, and last inspection date, and can be scheduled or executed immediately.

### C. Workflow

A typical scenario is planning an inspection and reviewing previous results. The flow of tasks and information is seen in Fig. 3. The process begins when an inspection is scheduled (monthly, weekly, or user-defined). The list of components for inspection is saved and on the scheduled date, the inspection workflow is triggered. The list is sent to the path planner to create a route for the robot and is then sent to the robot. At this point, the inspection process begins; the local navigation algorithm executes the route plan and performs any necessary corrections. After completing the inspection, the robot returns to the docking station for charging and performs a general status check. The collected data (images and metadata) are

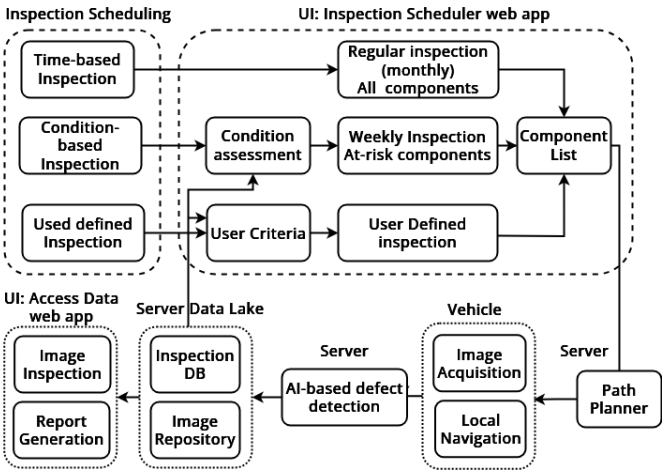


Fig. 3. DSS Workflow

uploaded to the server data lake. The database is updated with the metadata of each image: the unique identifier for the inspection associated with the images, the name of the component being inspected, the specific face of the component inspected, image timestamp, and image type (optical or thermal). The defect detection algorithm analyzes the images to assess wear levels. This process updates the optical and thermal degradation indices and the TRI for each component. A PDF report is generated for inspected components, including details such as name, images, TRI. The inspection process concludes with the new data available for user review along with access to historical inspection data.

## V. Inspection Planning

Given a set of components for inspection, the planner produces the optimal visitation order, ensuring safe and efficient navigation considering terrain visibility and total traversal cost. It operates in two phases: an offline preprocessing stage and an online planning stage. During preprocessing, local visibility maps are used to identify a set of highly visible waypoints per component. Those waypoints are organized into clusters using a greedy algorithm. For each visibility region, first we select the point with the highest visibility score and iteratively add new points at least a certain distance apart (e.g., 1.5m), up to a maximum of five or until a visibility threshold is reached. A fully connected graph is then constructed by computing paths between all such points across clusters, using A* on a costmap based on the traversability grid to assign edge weights. This costmap takes into account the robot footprint and adds an inflation around the objects to ensure safe navigation, assigning according costs. The extracted graph is stored as a weighted adjacency matrix for future use.

In the online phase, when a set of components are selected for inspection, the planner formulates a Generalized Traveling Salesman Problem (GTSP), where the goal is to visit exactly one point from each cluster to capture visual and thermal inspection data, while minimizing the total traversal cost. The goal is to solve the GTSP, where each cluster must be visited exactly once, by selecting one point from each cluster. To solve the problem We selected the GLNS solver [8], which is almost-optimal [9].

This solver follows the Adaptive Large Neighborhood Search (ALNS) framework [10]. It starts with a random initial tour and iteratively applies removal and insertion heuristics, whose weights are adaptively updated. At each step, a subset of vertices is removed and reinserted, with the new tour accepted or rejected using a simulated annealing criterion. The process includes an initial descent phase and several warm restarts with decreasing temperature. Various heuristics [11] and local optimization methods like Re-Opt and Move-Opt are used to refine solutions.

Since this problem is NP-hard, achieving guaranteed optimality is infeasible in the general case. However, for our use case, which ranges from 1 to 100 elements, after conducting extensive benchmarks our algorithm consistently yielded the optimal path.

## VI. Inspection Execution

### A. Autonomous Navigation

Autonomous navigation in EHVCs requires both reliable localization and robust handling of variable terrain conditions. To address these needs, we integrate (i) a global map for coarse path planning and (ii) a local real-time traversability module for continuous adaptation to newly observed obstacles and ground irregularities.

*1) Localization:* We utilize the localization pipeline introduced in [12] to maintain the robot pose on the 3D map. By coupling LiDAR scans with inertial measurements, the system provides stable pose estimates even in complex industrial terrains. All map-related updates (e.g., traversability information, costmap cells) are maintained in this global frame, thereby allowing the robot to navigate seamlessly between near-field hazard detection and broader inspection tasks.

*2) Real-Time Traversability Updates:* Although the offline-created global traversability map provides a foundational cost layer for path planning, the environment is subject to dynamic changes (e.g., newly placed equipment or unexpected debris). Therefore, a *local* traversability analysis is performed in real time based on [13]. This analysis supplements or overrides parts of the global costmap whenever new sensor data indicates impassable or risky terrain conditions. If the traversability exceeds a critical threshold, the corresponding cell is labeled as untraversable. Otherwise, the cell's danger value provides a local cost estimate that can be integrated into the global costmap, ensuring that newly observed obstacles or terrain hazards are quickly reflected in the robot's navigation decisions.

Our navigation setup is built on the standard `move_base` architecture in ROS, employing the **A*** algorithm as a *Global Planner* to compute a high-level path connecting two inspection points, while the local adjustments and real-time obstacle avoidance is handled the *Timed Elastic Band* planner.

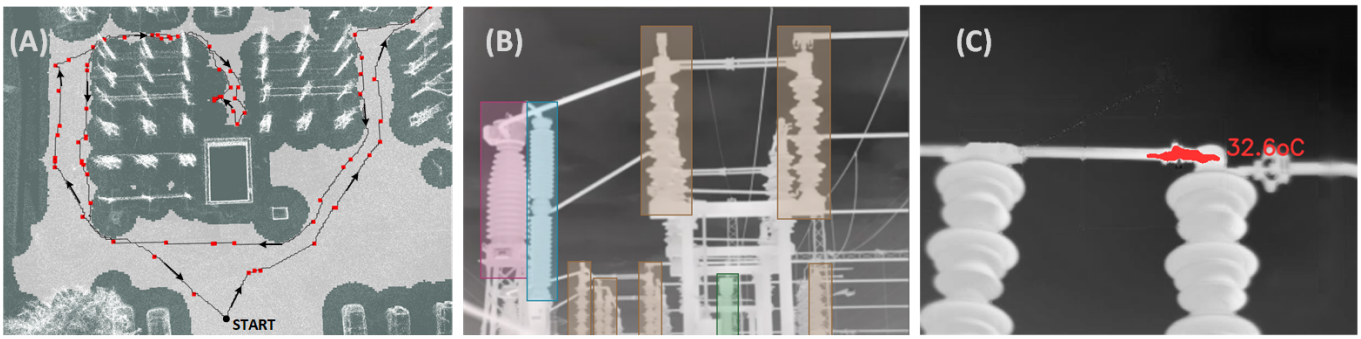


Fig. 4. A) Optimal path for an inspection plan. B) Bounding boxes of detected ECs in the thermal image. C) Detected fault.

### B. Defect Detection

In power substations, thermal imaging serves as the main inspection method, as early signs of damage in electrical components typically appear as uneven heat distribution across the equipment. Detecting components in thermal images is challenging due to the limited availability of large, annotated datasets and the lack of effective detection methods. To address this, we propose detecting ECs using the RGB channel and then identifying potential damage in the thermal image after registering it with the RGB image.

To detect ECs, we use YOLOv11n [14], a computationally efficient and high performance single-stage object detection method, suitable for deployment on mobile robots. We employ a pre-trained model and fine-tune it using manually annotated data from 30 minutes of video acquired from an actual power substation, in which 7 different classes of electrical components appear: surge arrester, bushings, current transformer, voltage transformer, isolator, circuit breaker and isolator (for the bottom part of all components). We annotated 1468 frames containing 7923 components in total. We employ a 80% - 20% split for fine-tuning and validation respectively, achieving 96.7% mean Average Precision and 94.0% F1 score.

To register detected bounding boxes to the thermal image, we follow a two-step process comprised of RIFT feature extraction, followed by RANSAC homography estimation (Fig. 4(B)). Potential faults are detected by computing connected components whose temperature is 10°C greater than the ambient temperature, approximated by the median temperature of all visible equipment (Fig. 4(C)). This rule-based approach is consistent with local guidelines on thermal inspection of power substations. We evaluated our algorithm on 84 annotated thermal images, extracted from reports composed by a professional inspector over a 2-year period, and applied our fault detection pipeline achieving 76/84 successful detections.

## VII. Conclusion

The ENORASI platform demonstrates a robust solution for automating inspections in EHVCs. By integrating advanced mapping, intelligent planning, real-time navigation, and defect detection, the system increases inspection frequency, reliability, and safety without increasing operational costs. Initial results confirm that it can effectively support condition-based maintenance strategies and improve the resilience of power infrastructure. Future work will focus on expanding system capabilities through AI-based predictive analytics and broader deployment in diverse substation environments.

## References

[1] Newsroom, "Extensive power cuts in attica affect transport," Aug 2018. [Online]. Available: https://www.ekathimerini.com/news/231924/extensive-power-cuts-in-attica-affect-transport/

[2] 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, 2022.

[3] 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, p. 501, 2016.

[4] P. Mermigkas, G. P. Moustris, C. S. Tzafestas, and P. Maragos, "Constructing visibility maps of optimal positions for robotic inspection in ultra-high voltage centers," in *33rd Mediterranean Conference on Control and Automation*, 2025.

[5] T. Flessa, G. Bonias, A. Skembris, D. I. Makrygiorgou, and D. Melissaris, "A framework for the robotic automation of ultra-high voltage center inspection tasks," *IET Conference Proceedings*, vol. 2024, pp. 219–224, 2025.

[6] R. Ahmad and S. Kamaruddin, "An overview of time-based and condition-based maintenance in industrial application," *Computers & Industrial Engineering*, vol. 63, no. 1, pp. 135–149, 2012.

[7] B. de Jonge, R. Teunter, and T. Tinga, "The influence of practical factors on the benefits of condition-based maintenance over time-based maintenance," *Reliability Engineering & System Safety*, vol. 158, pp. 21–30, 2017, special Sections : Reliability and Safety Certification of Software-Intensive Systems.

[8] S. L. Smith and F. Imeson, "GLNS: An effective large neighborhood search heuristic for the generalized traveling salesman problem," *Computers & Operations Research*, vol. 87, pp. 1–19, 2017.

[9] P. C. Pop, O. Cosma, C. Sabo, and C. P. Sitar, "A comprehensive survey on the generalized traveling salesman problem," *European Journal of Operational Research*, vol. 314, no. 3, pp. 819–835, 2024.

[10] S. Ropke and D. Pisinger, "An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows," *Transportation Science*, vol. 40, pp. 455–472, 11 2006.

[11] M. Fischetti, J. J. Salazar González, and P. Toth, "A branch-and-cut algorithm for the symmetric generalized traveling salesman problem," *Operations Research*, vol. 45, no. 3, pp. 378–394, 1997.

[12] 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 *2024 32nd Mediterranean Conference on Control and Automation (MED)*, 2024, pp. 107–112.

[13] T. Shan, J. Wang, B. Englot, and K. Doherty, "Bayesian generalized kernel inference for terrain traversability mapping," in *In Proceedings of the 2nd Annual Conference on Robot Learning*, 2018.

[14] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.