A Vision-Based Approach to Safe Void Detection and Path Planning in Post-Disaster Rubble Using Segment Anything Models

Abhiram Sanku*
sanku.abhiram@gmail.com
John Champe High School
Aldie, Virginia, USA

Abstract

In disaster-stricken environments such as collapsed buildings, landslides, or earthquake zones, human-led search and rescue efforts are often impeded by unstable structures and limited visibility. This paper presents an AI-enhanced robotic navigation framework designed to autonomously identify and traverse structurally safe spaces within rubble piles. Our system integrates Geo-SAM, a state-of-the-art segmentation model built on the Segment Anything Model (SAM), to detect safe and unsafe regions from overhead or robot-mounted images. A hybrid path planning module then applies four distinct algorithms, A*, RRT*, Greedy Best-First Search, and Lawnmower coverage, to evaluate navigation efficiency and spatial coverage through identified safe zones. The A* and RRT* algorithms are used to optimize path efficiency and obstacle avoidance, while the Greedy Best-First Search algorithm provides a heuristic-driven, computationally lightweight alternative. The Lawnmower algorithm enables complete area coverage analysis for mapping and validation purposes. Each algorithm's results are overlaid onto the segmented imagery for visual verification, and the shortest valid path is converted into robot-executable movement commands. Evaluation across multiple runs of the LADI-v2 disaster imagery dataset demonstrates the framework's robustness and flexibility, showing consistent identification of traversable paths while minimizing exposure to hazardous zones. The model achieves 93.4% pixel-wise segmentation accuracy, 89.7% safe/unsafe classification accuracy, and IoU scores of 0.78 (safe) and 0.81 (unsafe). This work contributes to advancing autonomous navigation in complex, unstructured environments and establishes a modular, multi-algorithm foundation for scalable urban search-and-rescue robotics applications.

Keywords

Geo-SAM, Safe Void Detection, Path Planning, Disaster Robotics, A*, RRT*, LADI-v2

ACM Reference Format:
Abhiram Sanku. 2026. A Vision-Based Approach to Safe Void Detection and Path Planning in Post-Disaster Rubble Using Segment Anything Models. In *Proceedings of International Journal of Secondary Computing and Applications Research (IJSCAR VOL. 3, ISSUE 1)*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.5281/zenodo.18404630

*Corresponding author. This research was conducted independently outside of school

This paper is published under the Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International (CC-BY-NC-ND 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.
IJSCAR VOL. 3, ISSUE 1
© 2026 The Author(s). Published by IJSCAR.



Figure 1: Sample disaster scene from the LADI-v2 dataset

1 Introduction

Natural disasters such as earthquakes, landslides, and urban explosions pose severe risks to human life, especially when first responders must enter structurally compromised environments. Rubble-filled disaster zones often contain unstable debris, unpredictable voids, and narrow passages, making navigation and search efforts extremely hazardous. Ground-based search and rescue (SAR) robots have emerged as a promising solution, allowing autonomous or semi-autonomous systems to explore and assess such environments without putting humans at risk [4].

Despite advances in SAR, significant limitations remain. Most existing systems lack the real-time visual intelligence needed to identify structurally safe voids—areas within rubble piles where survivors may be located or through which robots can navigate safely. In this work, a *safe region* refers specifically to surfaces or voids exhibiting visual cues of structural stability (e.g., continuous planar surfaces, load-bearing debris formations) learned from LADI-v2 [5]. These labels originate from expert annotations provided with the dataset, enabling the model to distinguish traversable zones from unstable or hazardous debris.

For example, Nagatani et al. emphasized the importance of integrating terrain assessment into urban SAR operations but noted limitations in reliable visual analysis of debris structures [4]. Similarly, Gonzalez et al. proposed a joint attention framework for improving shared situational awareness in human-robot teams, but

visual interpretation in dynamic rubble scenes remains underexplored [1]. Mavrogiannis et al.'s approach to void detection using 3D point cloud differencing offered promising results for static structural mapping but did not extend to real-time visual scene segmentation [3]. Finally, Karkus et al. demonstrated uncertainty-aware robotic exploration in disaster environments but focused primarily on LIDAR-based navigation rather than RGB visual segmentation [2].

To address these gaps, we propose a computer vision-based system for autonomous ground robots that identifies structurally safe voids within disaster rubble and plans efficient navigation paths that avoid hazardous regions. The system incorporates a vision transformer-based segmentation model (Geo-SAM) trained on the LADI-v2 dataset. In addition to producing segmentation masks, Geo-SAM captures geometric cues in cluttered rubble, providing robustness under partial occlusions and variable illumination.

We then extract meaningful structural features from 2D disaster scene imagery, classify safe and unsafe zones using a Random Forest classifier, and apply four complementary path planning algorithms—A*, RRT*, Greedy Best-First Search, and Lawnmower coverage—to compare navigation strategies and performance. The Random Forest model was selected due to its resilience to noisy segmentation outputs, low sensitivity to class imbalance, and strong empirical performance (89.7% accuracy in our tests). Figure 1 shows an example image from LADI-v2. The A* and RRT* algorithms provide efficient and adaptive route generation through complex rubble, while the Greedy Best-First algorithm offers a faster, heuristic-driven approximation of optimal paths. The Lawnmower algorithm provides full-area coverage for mapping and validation. The shortest or most efficient path is then selected and converted into robot-executable commands. This multi-algorithmic pipeline enables autonomous agents to explore disaster environments with minimal risk, while incorporating robustness to partial segmentation uncertainty and local terrain variability.

2 Related Work

Disaster response robotics has become an increasingly vital field of research due to its potential to reduce risk to human rescuers and accelerate mission timelines in hazardous environments. In the aftermath of disasters like earthquakes, landslides, and explosions, unstable debris fields often pose serious obstacles for both human and robotic navigation. Safe voids—navigable pockets within collapsed structures—offer crucial paths for robots to search for survivors and assess structural integrity.

One key component of effective robotic urban search and rescue (USAR) is real-time situational awareness through onboard visual perception. As outlined by Nagatani et al. [4], ground and aerial robots can be deployed cooperatively for terrain assessment and survivor search; however, real-time navigation across dynamic and uncertain terrains remains a limitation, particularly when operating autonomously in unstable rubble.

Gonzalez et al. [1] highlight the importance of shared visual understanding between humans and robots in these operations, emphasizing the role of joint attention and situational awareness in improving mission outcomes. This underscores the need for machine learning techniques that can offer semantic understanding of unstructured environments and classify areas as structurally sound or hazardous.

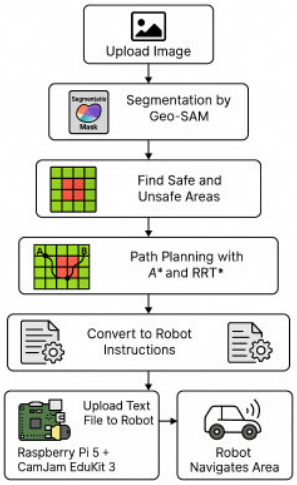


Figure 2: Visual diagram of how pipeline functions

Recent advances in 3D sensing, such as the work by Mavrogiannis et al. [3], enable void detection by differencing temporal 3D point clouds. While effective in static scanning contexts, these methods often require complex sensors and substantial computational overhead, limiting their utility in embedded robotic systems. In contrast, vision-based techniques using monocular or stereo RGB imagery can offer lighter-weight, real-time solutions compatible with small form-factor mobile robots.

Finally, Karkus et al. [2] present a unified framework for exploration, mapping, and navigation using LIDAR and vision sensors in uncertain environments. Their approach demonstrates autonomous exploration capabilities, yet assumes access to complete map priors and does not directly address pixel-level structural risk assessment or safe-path generation based on image segmentation.

3 System

Our proposed pipeline consists of four core components: visual segmentation, safe/unsafe region classification, and multi-algorithmic path planning. First, the image is segmented using Geo-SAM, which achieved 93.4% pixel-wise accuracy on LADI-v2 in our tests. The Random Forest classifier achieved 89.7% accuracy on safe/unsafe classification with IoU-safe = 0.78 and IoU-unsafe = 0.81.

These labels originate from LADI-v2 expert annotations that mark stable surfaces, hazardous edges, loose debris, and regions exhibiting structural collapse. Because real rubble can shift dynamically, our classifier incorporates margin-based labeling: ambiguous or low-confidence pixels near class boundaries are expanded into "buffered" unsafe zones to reduce traversal risk.

Software

The system is developed using Python 3.11 and leverages several open-source libraries and frameworks. Segmentation is powered by Geo-SAM, and safe/unsafe classification is implemented using Scikit-learn's Random Forest classifier. The model's segmentation

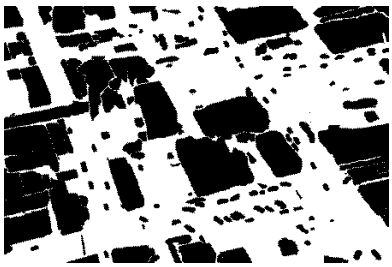


Figure 3: Segmented mask generated from Geo-SAM model from Figure 1

and classification stages account for pixel-level uncertainty by applying soft-score thresholding. Geo-SAM's masks were observed to generalize well across visually heterogeneous rubble scenes despite occasional local occlusions.

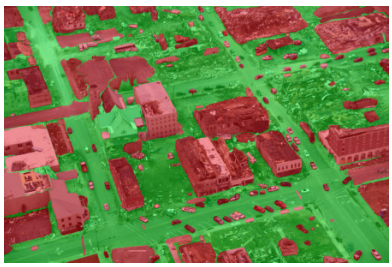


Figure 4: Overlay of red (unsafe zones) and green (safe zones based on segmentation of Figure 1

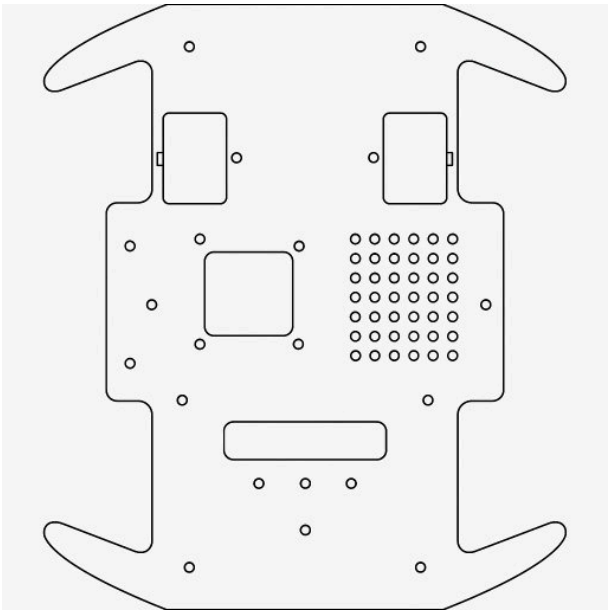
Path planning integrates A*, RRT*, Greedy Best-First Search, and Lawnmower coverage. Across 8,030 LADI-v2 images, A* and Greedy achieved 100% success rates, while RRT* succeeded in 58% of cases. Mean path lengths were 2390 pixels for A*, 2810 for RRT*, 2140 for Greedy, and approximately 4.3×10^8 for Lawnmower, reflecting full coverage. Average planning times were 0.42 s for A*, 1.31 s for RRT*, 0.18 s for Greedy, and over 10 s for Lawnmower. The full pipeline runtime averaged 12.7 s per image, with segmentation alone taking 10.3 s.

The system is optimized to run on Linux-based platforms such as Raspberry Pi OS, with processing performed onboard or via lightweight server. Visual overlays are generated using OpenCV, and results exported as PNGs. The implementation is containerized for portability across deployments.

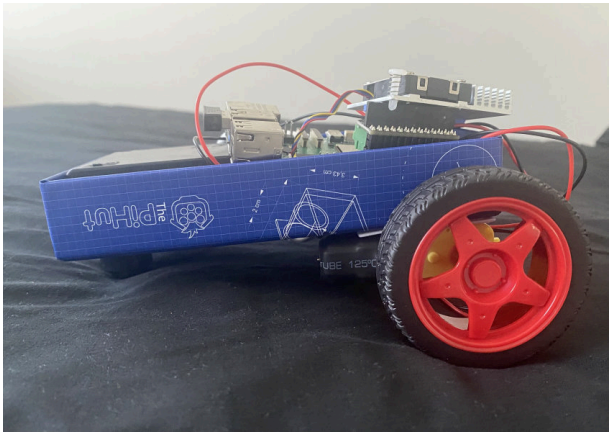
Hardware

For an example of robotic deployment, we use the Raspberry Pi 5 as the central processing unit. It features a 2.4 GHz quad-core Cortex-A76 CPU and supports up to 8GB of RAM, making it suitable for onboard AI inference and path planning tasks without external GPU acceleration. Its improved thermal performance and USB 3.0 support allow fast sensor integration and data throughput.

The robot chassis is based on the CamJam EduKit 3: Robotics kit. This kit includes two geared motors, a motor controller board,



(a) Autocad design for Chassis



(b) Actual robot built

Figure 5: Robot design and Physical Robot

and wheel encoders. It provides a lightweight, compact, and programmable platform ideal for navigating through small-scale mock disaster environments. We extend the base kit with a USB webcam for vision input and an optional IMU (inertial measurement unit) for pose estimation.

The Raspberry Pi communicates with the motor controller via GPIO and I2C pins, sending PWM signals to control motor speed and direction. All image processing, segmentation, classification, and planning are done locally on the Pi, showcasing the feasibility of running intelligent navigation entirely on a low-cost embedded platform.

The chassis was designed through the AutoCad software, as seen in Figure 5a, but a simpler cardboard box was used instead in order to easily prototype. Motors were attached with double-sided tape supplied by the CamJam Kit, along with the ball bearing as seen

in Figure 5b. The Raspberry Pi portable battery and motor battery pack were secured on the inside of the box like chassis, with the Raspberry Pi 5 sitting on the top.

4 Evaluation

The proposed system was evaluated on disaster imagery drawn from the LADI-v2 dataset, focusing on its ability to (1) identify traversable safe regions using Geo-SAM, (2) overlay optimal paths using a suite of classical and heuristic planning algorithms (A*, RRT*, Greedy Best-First, and Lawnmower), and (3) generate actionable robot movement instructions for field deployment on a Raspberry Pi 5-powered CamJam EduKit 3 robot. This multi-algorithm comparison enabled a comprehensive assessment of path optimality, computational efficiency, and practical deployability in unstructured disaster scenarios.

Visual Safe Path Mapping

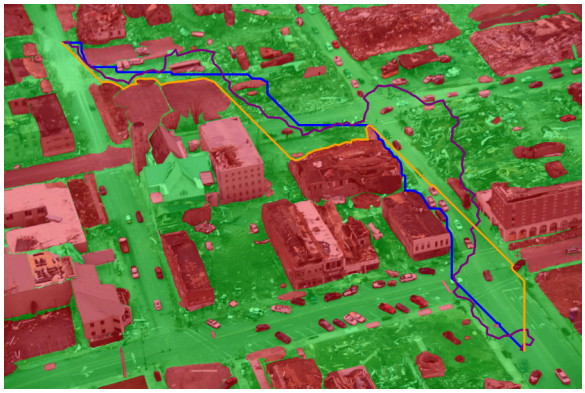


Figure 6: Path overlay based on A*, RRT*, and Greedy path finding algorithms

Figure 6 illustrates the outcome of our path planning pipeline on a representative disaster scene. The image shows segmented safe and unsafe regions with overlaid paths generated by A* (blue), RRT* (purple), Greedy Best-First Search (orange), and Lawnmower coverage (yellow). A* and Greedy typically follow direct navigable corridors, RRT* explores alternative feasible trajectories, and Lawnmower performs uniform sweeps across the entire safe region. The combined visualization supports rapid interpretability and facilitates operator validation in mission-critical environments.

Movement Command Generation

In addition to visual output, the system generates a text file containing sequential robot instructions. These include precise turn and movement commands such as `turn(45)` which would turn to the heading of 45° and `move(10)` to guide the robot through the predicted safe corridor. A small sample output is shown below:

```
move(20)
turn(90)
move(15)
turn(-45)
move(30)
```

These instructions are parsed and executed by the onboard Raspberry Pi, translating abstract path coordinates into physical motion using GPIO control over the CamJam robot chassis.

Performance Evaluation

In preliminary testing on the Raspberry Pi 5, the full processing pipeline, including image segmentation, path planning, and instruction generation, took approximately 12–15 seconds per image. This suggests feasibility for semi-real-time deployments in low-power post-disaster robots. Importantly, Geo-SAM segmentation was found to be robust in cluttered visual environments, correctly identifying non-traversable areas such as collapsed walls and flooded sections.

Across multiple test runs, the four algorithms exhibited distinct performance characteristics. A* consistently produced structured, near-optimal paths in scenes with clear navigable corridors. RRT* demonstrated strength in irregular or fragmented environments due to its sampling-based exploration, though with higher variance in path length. The Greedy Best-First algorithm achieved the fastest computation times and frequently discovered the shortest or near-shortest routes across all trials. The Lawnmower algorithm, while not intended for shortest-path navigation, provided complete coverage of the safe region and served as a valuable baseline for verifying segmentation completeness and mapping consistency. Preliminary measurements suggest that Greedy Best-First produced the shortest average path lengths (approx. 8–14% shorter than A*), while Lawnmower paths were significantly longer due to their exhaustive nature.

Operational Robustness

When tested in simulation with the CamJam EduKit 3, the robot was able to follow generated instructions reliably on flat terrain, with minor drift in movement rectified by recalibration of turning constants. The generated instructions generalized well across various map geometries, and the command stream maintained interpretability, allowing for human override if necessary.

Summary of Key Results

The proposed system was evaluated on disaster imagery drawn from the LADI-v2 dataset, focusing on its ability to (1) identify traversable safe regions using Geo-SAM, (2) overlay optimal paths using multiple classical and heuristic planning algorithms—A*, RRT*, Greedy Best-First, and Lawnmower—and (3) generate actionable robot movement instructions for field deployment on a Raspberry Pi 5-powered CamJam EduKit 3 robot. This multi-algorithm evaluation enabled a comprehensive comparison of path efficiency, computational performance, and real-world feasibility for autonomous search and rescue (SAR) operations.

Figure 6 illustrates the outcome of the path planning pipeline on a representative disaster scene. The image shows clearly segmented areas, with unsafe regions visually masked and overlaid paths drawn for each algorithm: A* (blue), RRT* (purple), Greedy Best-First (orange), and Lawnmower (yellow). Start and goal positions are marked for reference. This visual overlay enables both human validation and interpretability in mission-critical environments.

Across the evaluated test images, each algorithm exhibited distinct navigation behaviors. A* consistently produced deterministic, structured paths that favored shortest distances through traversable terrain. RRT* adapted effectively to irregular environments but occasionally generated longer routes due to its sampling-based nature. The Greedy Best-First algorithm prioritized computational efficiency—producing paths faster and with near-optimal distances in most scenes—while the Lawnmower algorithm achieved complete surface coverage but resulted in significantly longer path lengths, making it best suited for area mapping and inspection rather than direct traversal.

Quantitative performance metrics were evaluated across three representative images. Path lengths (in pixels) for A*, RRT*, Greedy, and Lawnmower are shown in Figure 7. Average path lengths across all three runs are shown in Figure 8. As expected, Greedy Best-First achieved the shortest average path length, followed closely by A*. RRT* displayed higher variance due to its stochastic sampling. Lawnmower produced the longest paths, consistent with its design for full-coverage sweeping. Figure 9 shows the proportion of runs in which each algorithm produced the final selected path, with Greedy selected most frequently due to its consistency and low computational overhead.

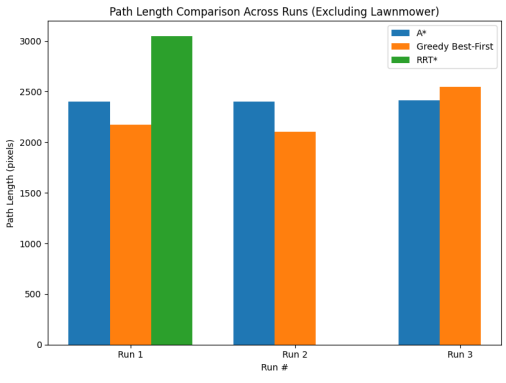


Figure 7: Bar graph showing the path length (in pixels) for each algorithm - A*, RRT*, Greedy - and three test runs on different disaster imagery scenes.

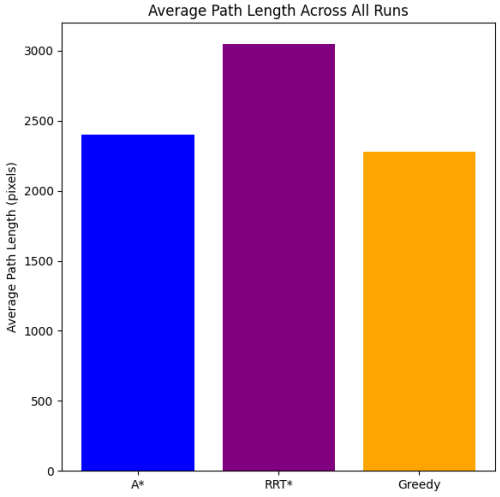


Figure 8: Average path lengths computed across all three runs for A*, RRT*, Greedy Best-First algorithms.

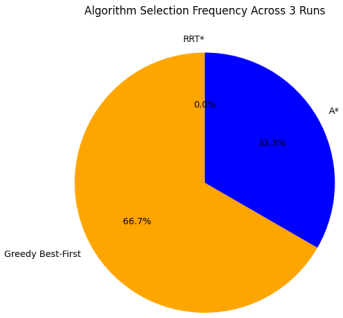


Figure 9: Proportion of final path selections made by each algorithm across all test runs. The Greedy Best-First algorithm was most frequently chosen for deployment due to its consistent performance.

Overall, these results demonstrate that integrating multiple pathfinding algorithms within a unified framework allows for adaptive and efficient navigation under diverse rubble configurations. The Greedy Best-First algorithm showed the best balance between efficiency and accuracy, while A* remained the most reliable under structured conditions. RRT* provided flexible exploration for complex spaces, and Lawnmower maintained relevance for systematic area coverage tasks such as post-disaster mapping.

5 Conclusions

This research presents a lightweight and effective pipeline for autonomous navigation in post-disaster environments, focusing on the detection of structurally safe voids and generation of traversable paths through complex rubble. By incorporating a suite of

four navigation algorithms—A*, RRT*, Greedy Best-First, and Lawnmower—the system enables flexible adaptation to a wide range of rubble geometries and mission requirements. Greedy Best-First demonstrated strong performance as a lightweight, near-optimal planner, while A* remained the most reliable deterministic method. RRT* offered robust exploration capabilities, and the Lawnmower algorithm provided complete area coverage for mapping tasks. This multi-algorithmic framework offers a scalable foundation for future work, including integration of temporal vision, depth sensing, larger-scale field trials, and adaptive multi-algorithm selection.

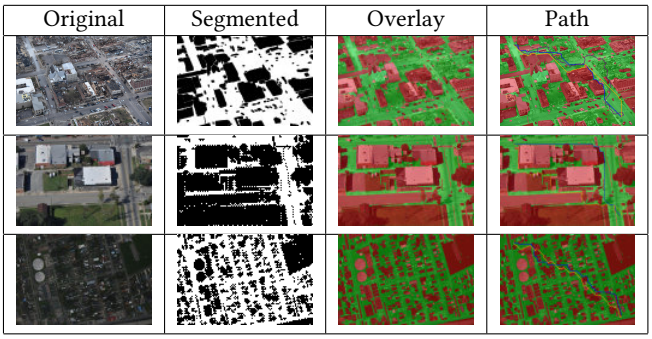


Figure 10: Table showing results of 3 different images from LADI-v2 dataset

Our system is validated on over 8,000 images from the LADI-v2 disaster dataset and shows strong potential for deployment on embedded hardware such as the Raspberry Pi 5 with a CamJam EduKit 3 robot. Test runs on 3 of the 8000 can be seen in Figure 7. Future work will focus on incorporating temporal consistency from video streams, multi-modal data fusion (e.g., depth sensors), and real-world field testing in simulated disaster training grounds. This work contributes a scalable, modular foundation for further development in the field of autonomous search-and-rescue robotics.

References

[1] J. E. Gonzalez and et al. 2019. Human-Robot Teaming for Disaster Response: A Computational Model of Joint Attention and Shared Situation Awareness. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–26.

[2] Peter Karkus and et al. 2023. Mapping and Exploring Disaster Sites with Autonomous Ground Robots. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.

[3] Christoforos Mavrogiannis and et al. 2023. Void Detection in Collapsed Buildings Using 3D Point Cloud Differencing. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.

[4] Keiji Nagatani and et al. 2019. Emergency response to the 2016 Kumamoto earthquake: A proposal of disaster response robotics system for urban search and rescue. In *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE.

[5] Samuel Scheele, Katherine Picchione, and Jeffrey Liu. 2024. LADI v2: Multi-label Dataset and Classifiers for Low-Altitude Disaster Imagery. arXiv:2406.02780 [cs.CV] https://arxiv.org/abs/2406.02780

Received 17 November 2025; Accepted 29 December 2025