



ARL-TN-1171 • SEP 2023



A Barrier Detection Experiment

by Craig Lennon, Luis Hernandez, Seth Ellis, Andre Harrison,
and Marshal Childers

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.



A Barrier Detection Experiment

Craig Lennon, Luis Hernandez, Andre Harrison, and Marshal Childers
DEVCOM Army Research Laboratory

Seth Ellis
Oak Ridge Associated Universities

REPORT DOCUMENTATION PAGE

1. REPORT DATE		2. REPORT TYPE		3. DATES COVERED	
September 2023		Technical Note		START DATE 9 January 2023	END DATE 27 January 2023
4. TITLE AND SUBTITLE A Barrier Detection Experiment					
5a. CONTRACT NUMBER		5b. GRANT NUMBER		5c. PROGRAM ELEMENT NUMBER	
5d. PROJECT NUMBER		5e. TASK NUMBER		5f. WORK UNIT NUMBER	
6. AUTHOR(S) Craig Lennon, Luis Hernandez, Seth Ellis, Andre Harrison, and Marshal Childers					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) DEVCOM Army Research Laboratory ATTN: FCDD-RLA-JC Aberdeen Proving Ground, MD 21005				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TN-1171	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.					
13. SUPPLEMENTARY NOTES ORCID IDs: Andre Harrison, 0000-0002-6850-9677; Marshal Childers, 0000-0002-8723-6977					
14. ABSTRACT In January 2023, an assessment of research into the use of computer vision for barricade detection was conducted by the Artificial Intelligence for Maneuver and Mobility Essential Research Program. The computer vision algorithm was intended to identify routes that would be traversable except for the presence of movable objects blocking movement along the route. The experiment described here was conducted to assess how the algorithm performed at detecting a variety of barricades across a range of distances and angles. The algorithm, as integrated onto a robotic platform, detected barricades at a distance of approximately 10–15 m. The detection was more reliable when the robot was directly facing the barricade than when the robot approached at an angle. The difference in reliability could have been due either to the algorithm or to a camera field of view that did not include the barricade when the robot entered the 15-m detection range at an angle.					
15. SUBJECT TERMS Military Information Sciences, computer vision, autonomy, robotics, artificial intelligence					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED	UU	31	
19a. NAME OF RESPONSIBLE PERSON Craig Lennon				19b. PHONE NUMBER (Include area code) (410) 278-9886	

STANDARD FORM 298 (REV. 5/2020)

Prescribed by ANSI Std. Z39.18

Contents

List of Figures	iv
List of Tables	iv
1. Introduction	1
2. Robotic System and Algorithm	2
3. Experimental Methods	4
4. Results	10
5. Conclusion	16
6. References	17
Appendix. Complete Record of Results	18
List of Symbols, Abbreviations, and Acronyms	24
Distribution List	25

List of Figures

Fig. 1	Example barrier scenario	1
Fig. 2	Warthog robotic platform	2
Fig. 3	An image (left) with segmentation (center) and detection (right)	3
Fig. 4	An example of semantic segmentation from the RELIS-3D data set (Jiang et al. 2021).....	3
Fig. 5	Experimentation site image from above	5
Fig. 6	Legend for Figs. 7–12	6
Fig. 7	Avenues of approach for runs 1 and 2, and the vehicle barrier	7
Fig. 8	Avenues of approach for runs 3–12	7
Fig. 9	Barrier 1 (left), barrier 2 (center), and barrier 3 (foreground, right).....	8
Fig. 10	Avenues of approach for runs 13–16	8
Fig. 11	Avenues of approach for runs 16–20	9
Fig. 12	Avenues of approach for runs 21–24	9
Fig. 13	Logs.....	10
Fig. 14	An image, segmentation, and barrier detection from 10 m.....	12
Fig. 15	A 45-m detection with a false positive	12
Fig. 16	Detection distances by run and approach angle.....	13
Fig. 17	Max correct and min incorrect detection	14

List of Tables

Table 1	Results for run 1	11
Table 2	Analysis of variance for 19 runs, including nondetections.....	15
Table 3	Analysis of variance for 13 runs, omitting nondetections	16
Table 4	Analysis of variance for 19 runs, including nondetections.....	16
Table A-1	Complete list of runs and results for each run	20

1. Introduction

In January 2023, the US Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory (ARL) Artificial Intelligence for Maneuver and Mobility Essential Research Program (AIMM ERP) conducted experiments at the ARL Robotics Research Collaboration Campus (R2C2) at Graces Quarters, Maryland. The goal of these experiments was to assess the ability of a barrier detection algorithm to identify a set of obstacles as a barrier. Here, a barrier is defined as an object, or set of objects, that split an otherwise continuous, traversable region into two disjoint regions. The barrier detection research effort is part of AIMM's Context-Aware Decision-Making research thrust. The barrier detection algorithm takes as an input the semantic class and relative location of objects in the surrounding environment and decides whether the objects have the functional role of a barrier. Functional role, in this report, is defined as the relations between the physical properties of an object, its placement in the environment, and a specific goal (Hodges 1995). As an example, a vehicle may be parked in front of a building, in the middle of the field, or blocking a path between two buildings. In each case the object remains the same, but in the last case its relation to other objects within the environment modifies its role within the environment, potentially changing how the object is treated within a given mission, and whether the vehicle should be classified as a barrier (Fig. 1). The physical property is the size of the vehicle, and its placement in a position blocking otherwise traversable terrain means it may be there to prevent the goal of moving through that terrain. This combination of physical size, placement, and the fact that its size and placement would prevent a potential goal result in the vehicle being identified as a barricade. In subsequent sections of this report, we describe: 1) the barrier detection algorithm and its place as part of the robotic system used in experimentation, 2) the experimental procedures, 3) an analysis of the results, and 4) conclusions.



Fig. 1 Example barrier scenario

2. Robotic System and Algorithm

The robotic platform used in the experiment was a Clearpath Warthog, shown at the test site in Fig. 2. Throughout the experiment, the robot movement was controlled via teleoperation. The robot was connected to an external laptop through which researchers could monitor it, but it used onboard sensing and computation for mapping its environment and for detecting barriers. For the terrain mapping, the robot used a laser imaging, detection, and ranging (LiDAR) sensor and a front-mounted camera to create a sparse depth map and a dense semantic segmentation map of the terrain. In this process, an image from the camera has a semantic label applied to each pixel, and the set of LiDAR points that fall within each semantic region are then used to provide a rough depth estimate for each segment in the image. An example of this process used for barrier detection is shown in Fig. 3, which shows the barrier detection algorithm applied to images from the Deep Green challenge test results (Deep Green 2023). Deep Green was an Army competition, held in 2022–2023, in which participants competed to train a semantic segmentation algorithm to achieve the best possible accuracy on a data set selected by the Army Office of Business Transformation and ARL, which co-sponsored the competition.



Fig. 2 Warthog robotic platform




Fig. 3 An image (left) with segmentation (center) and detection (right)

An example of semantic segmentation (Ellis and Harrison 2023) is shown in Fig. 4, which shows scenes with labeled terrain (e.g., grass) and objects (e.g., barrier).




Fig. 4 An example of semantic segmentation from the RELIS-3D data set (Jiang et al. 2021)

Ellis and Harrison (2023) provide a detailed description of the barrier detection algorithm, but we provide a brief summary of how it works here. The detection algorithm uses multimodal sensor data to discriminate between barriers and obstacles within the current camera frame, and then labels the objects that make up a barrier based on if they are considered movable or environmental elements. The algorithm takes as input a semantic segmentation map of the red, green, blue (RGB) image, and an ontology in which classes are described in terms of difficulty of traversal. To categorize regions as traversable or non-traversable, the algorithm takes class labels from the semantic segmentation and, based on the ontology, separates the classes into traversable or non-traversable depending on the difficulty of traversal. This division of the classes in the image results in binary masks that represent the regions of traversal and the regions of non-traversal, respectively.

Based on the layout of traversable and non-traversable regions, the algorithm employs one of two approaches to identify if any of the non-traversable regions is a barrier or just an obstacle. One approach applies to images in which there are two or more traversable regions in the image frame, and the second approach is used when there is only a single traversable region in the image frame.

The first (multiregion) approach leverages the idea that if there are multiple visible regions of traversal contained within the scene, then the shortest distance between the distinct regions is where the potential barrier is located. Barriers are modeled as horizontal blockades that separate traversable regions, creating two or more disjoint regions. Once the barrier is located, the algorithm:

- 1) Identifies the original semantic class of the non-traversable region (barrier).
- 2) Categorizes the barrier as environmental (immovable) or movable based on the class labels from step 1, and the height of the barrier. Barriers greater than a certain height (set as a parameter in the algorithm) are considered environmental.
- 3) Connects and identifies as a barrier the multiple classes/regions if the non-traversable region consists of more than one class, or if multiple non-traversable regions are in close enough proximity to block the platform.

The multiregion approach would not apply to a situation in which non-traversable regions at different depths, like a series of offset barriers at a checkpoint, block movement enough to constitute a barrier but are continuous in the image. For the case in which there is one region of traversal, the algorithm compares the width and depth of traversable and non-traversable regions in a 2.5-dimensional space, which is defined as a space in which surfaces need not be flat, as in 2-D, but in which volume is not considered. If the range of depths of the non-traversable region is within the range of depths of the single, traversable region, and the non-traversable region is sufficiently wide, it is considered a barrier. The non-traversable region is considered a barrier if the non-traversable region spans 90% of the scene, with the 90% allowing for the gaps between objects that might be present in the single-region approach.

3. Experimental Methods

The experimentation took place at the R2C2 experimentation facility. A satellite image from Google Maps (2023) is shown in Fig. 5 in which the buildings shown on the image are overlaid with blue rectangles. Subsequent figures of the site will consist of diagrams showing the rectangles only, without the underlying imagery, which will allow the details relevant to the experiment to be shown more clearly.

A yellow star is placed on the same location in each diagram as a reference point to allow the reader to understand the orientation with respect to the image in Fig. 5.




Fig. 5 Experimentation site image from above

There were a total of 20 runs conducted during the experiment, each from a distinct combination of position, angle, and obstacle type. This allowed greater variety in experimental conditions, but at the sacrifice of understanding how the algorithm might vary during replicated runs. This choice was made since the researchers were confident that variation between repeated runs would be low, but variation across locations, angle of approach, and type of obstacle could be substantial, and so was of greater interest. The type of obstacles used in the barriers are shown in Fig. 6, which is also the legend for Figs. 7, 8, and 10–12. The barriers include the vehicle shown in Fig. 7 as well as four barriers described as follows:

- Barrier 1: orange transit cones with a mannequin (Fig. 9)
- Barrier 2: a line of yellow cones (Fig. 9)
- Barrier 3: a line of pallets with transit cases on top of some (Fig. 9)
- Barrier 4: a pile of logs (Fig. 13)

Diagrams showing the avenues of approach and images of the obstacles are shown in Figs. 7–12.




Fig. 6 Legend for Figs. 7–12

In each of the 24 runs, the robot was teleoperated. For all runs except 9–12, the robot stopped at set distances from a single barrier. The algorithm’s classification of the object as a barrier, along with imagery and LiDAR data, was saved for later analysis. Imagery from the RGB camera sensors were recorded at 30 fps and LiDAR point clouds were recorded at 10 Hz. At each of the locations where the robot stopped, the algorithm was evaluated by a researcher as to whether:

- 1) The semantic segmentation algorithm appropriately divided the terrain and objects into well-defined regions, which would allow terrain to be recognized as traversable and for objects to be recognized as blocking the traversable terrain.
- 2) The semantic segmentation algorithm labeled the regions with a correct label.
- 3) The barrier detection algorithm detected the object as a barrier.
- 4) The detection made in (3) was correct given the segmentation and labeling in (1) and (2), in the opinion of the researcher.

For runs 9–12, the vehicle was driving through the site past multiple obstacles and did not stop, nor were assessments of barrier detection made at varying distances. These runs can be considered as part of a data collection rather than an evaluation of the algorithm. In the following descriptions and figures, we show the avenues of approach as well as the location and types of obstacles, and list distances at which the robot stopped.

Runs 1 and 2 started near the R2C2 administrative facilities and approached a vehicle parked between two buildings. The runs started either from a position that could see the vehicle the entire time (run 1), or a position that could not (run 2). The robot stopped at distances of 60, 45, 30, 15, 10, and 5 m for each run.




Fig. 7 Avenues of approach for runs 1 and 2, and the vehicle barrier

Figure 8 shows the directions of approach for runs 3–12. Runs 3 and 4 approached the same vehicle as runs 1 and 2, but from a different angle. Runs 5 and 6 approached barrier 1, while runs 7 and 8 approached barrier 2. The barriers are shown in Fig. 9. For runs 3–8, the robot stopped at distances of 60, 45, 30, 15, 10, and 5 m. For runs 9–12, the robot drove through or around the site to collect data.




Fig. 8 Avenues of approach for runs 3–12



Fig. 9 Barrier 1 (left), barrier 2 (center), and barrier 3 (foreground, right)

Runs 13–16 approached barrier 3 (Fig. 12) from different directions, as shown in Fig. 10, stopping at distances of 60, 45, 30, 15, 10, and 5 m from barrier 3.




Fig. 10 Avenues of approach for runs 13–16

After run 16, barrier 3 was modified so the robot could drive past it while approaching barrier 2 for run 17. The robot approached barrier 2 again in run 18, and barrier 1 in runs 19 and 20. In each case the barrier detection algorithm was assessed at distances of 60, 45, 30, 15, 10, and 5 m.




Fig. 11 Avenues of approach for runs 16–20

For runs 21–24, the barrier was a pile of logs, as shown in Fig. 13. A diagram of the site, alongside an overhead satellite image from Google Maps (2023) is shown in Fig. 12. In each case, the barrier detection algorithm was assessed at distances of 60, 45, 30, 15, 10, and 5 m. With run 21, technical difficulties interfered with the data collection, so the resulting data is not analyzed.



Fig. 12 Avenues of approach for runs 21–24



Fig. 13 Logs

In Section 4, we examine the results from runs 1–8, 13–20, and 22–24, with the goal of understanding the effect of distance and angle of approach on the barrier detection algorithm. Runs 9–12, in which the vehicle was guided through the whole facility, were examined by the researchers, but were not structured the same as the other runs and are not analyzed in Section 4.

4. Results

The goals of the experiment were first to ensure the algorithm would function as part of a robotic platform, and second to understand the relation of distance and angle to barrier detection and how detection varied across different combinations of location and object. The algorithm did function on the platform, meeting the first goal. To achieve the second goal, we examined the performance of the algorithm in terms of the distance at which the barrier was detected, and the relation of that distance to angle and combinations of locations and object. The results of all runs are included in the Appendix, and an example result is shown in Table 1.

These are the variables we will be considering in our analysis:

- **Detected** is a binary variable (yes/no) that indicates whether the obstacle was detected as a barrier at a given distance. In Table 1, detected is “yes” at 45 m, as indicated by the “YES” in the Detected Barrier column at 45 m. Only YES/NO results are analyzed, although for some runs there are other entries, such as NA or “out of range,” which indicate that the relative position of the robot to the barrier, or something else related to the experiment, prevented the barrier detection algorithm from operating.
- **Max Detected** is the greatest distance in meters at which the barrier was detected, for example, 45 m in Table 1.

- **Final Detected** is the greatest distance in meters such that there were no “NO’s” in Detected Barrier at any closer distance. For example, 15 m in the table (run 1). This is the greatest distance at which one could feel confident that the barrier would be detected.
- **Pose** indicates the angle at which the robot approached the barrier and could be either close to perpendicular (P) or at an oblique angle (O).
- **Site** indicates the combination of barrier and location. Since each barrier was used at only one location, site is referenced by the barrier type: vehicle, barrier 1, 2, 3, or logs.
- **Max Correct Detection** is the maximum distance at which the barrier is detected, and the researcher said it was the correct calculation. It is the best true positive.
- **Min Incorrect Detection** is the minimum distance at which the barrier detection algorithm detected the barrier, but the researcher believes the detection to be incorrect based on the semantic segmentation. It is the worst false positive.

Table 1 Results for run 1

Run	Distance (m)	Pose	Well-defined regions	Regions are consistent	Detected barrier	Correct calculation
1	60	P	YES	YES	NO	YES
1	45	P	YES	YES	YES	NO
1	30	P	YES	YES	NO	YES
1	15	P	YES	YES	YES	YES
1	10	P	YES	YES	YES	YES
1	5	P	NA	NA	NA	NA

We begin by plotting Max Detected and Final Detected to see how they vary by pose and by site. In runs 6 and 22, the barrier was never detected, so for those runs, Max Detected and Final Detected are set to zero, while in runs 23 and 24 the barrier was detected early on, then not detected, so only Final Detected is set to zero. Max Detected must be at least as great as Final Detected because of how the variables are defined, and Max Distance must be at least as great as Max Correct Distance. Figures 14 and 15 show examples of images and segmentation leading to correct (Fig. 14) and incorrect (Fig. 15) barrier detection. In Fig. 14, the barrier is correctly detected from a range of 10 m. Figure 15 is an example of a false positive, in which the lack of resolution, due to distance, resulted in incorrect detection. A vehicle was present, but because no drivable region was identified around or beyond the vehicle, it should not have been identified as a barricade. Boxes in Fig. 15 show the vehicle (left), which is correctly segmented as such (right). However, because there is no

drivable region (green) in the center image, the vehicle should not have been shown as a barricade (pink) in the center image. An inspection of Fig. 16 shows that it is reasonable to expect detection within 10–15 m of the barrier, although a few 45-m detections did occur. In the opinion of the researchers, however, these 45-m detections may be detections made without sufficient reason for a detection. For each 45-m detection, researchers recorded the value of Correct Calculation as “No.” In Fig. 16, red indicates a perpendicular value of Pose, and black an oblique value of Pose. Circles represent Max Detected distance while squares represent Final Detected distance. When the squares and circles coincide, this indicates a consistency in the algorithm’s classification, while circles above squares represent runs on which the algorithm detected the barricade (perhaps incorrectly), then did not detect it at a closer range. The average distance of Max Detected and Final Detected are 22.6 and 12.4 m, respectively.




Fig. 14 An image, segmentation, and barrier detection from 10 m




Fig. 15 A 45-m detection with a false positive




Fig. 16 Detection distances by run and approach angle

Consequently, we examine in Fig. 17 the maximum distance at which the barrier was detected correctly (in the opinion of the researchers), and also consider the minimum distance at which the barrier was not detected, incorrectly (in the opinion of the researchers). The triangles in Fig. 17 show Max Correct Distance, the maximum distance at which researchers judged the detection to be correct. No correct detection was made in runs 4, 8, 17, 22, 23, and 24, so those runs were assigned a value of 0. Excepting run 15, correct detections occur at approximately 15 m from the barrier. The diamonds in Fig. 17 show the minimum distance at which researchers judged an incorrect detection to have occurred. The plot from Fig. 17 (plot correct) suggests that 15 m is the present limit for detection, but that detection at 15 m is inconsistent.

A comparison of red versus black triangles also suggests that detection occurs from farther away when the angle is close to perpendicular. However, this may be because so many runs at oblique runs never detected the barrier, so have values of zero. This evidence is difficult to interpret, because the camera's field of view is limited, meaning that when a barrier is approached at an oblique angle, it may have left the field of view before being detected. It is possible that given the range at which the algorithm correctly detects barriers, detection of a barrier from an oblique angle is not possible with the camera's field of view. This would be a limitation of the robot and the algorithm when considered together as a system. It is also worth noting that the range of detection is determined by the camera image and might be

increased by a camera that zooms in on, or otherwise focuses on, the region containing the barrier.

In Fig. 17, red indicates a perpendicular value of Pose, and black an oblique value of Pose. Triangles represent Max Correct Detection distance while diamonds represent Min Incorrect Detection distance. The average Max Correct Detection distance is 15.4 m.




Fig. 17 Max correct and min incorrect detection

To better understand the relationship of detection distance to Pose, we consider two possibilities:

- A) The angle of approach (Pose) is relevant because there are so many failed detections with the oblique angle of approach.
- B) The angle of approach (Pose) is relevant even without considering the failed detections.

We find (A) supported by comparing the results of analysis of variance with and without the failed detections included in the data set. Both analyses, shown in Tables 2 and 3, consider response as the maximum distance at which the barrier is correctly detected (Distance) with the angle of approach (Pose) and the type of barrier (Site) as factors. The Max Correct Detection distance for Pose i at Site j is modeled as an overall average mean max detection distance μ with changes to that distance caused by the value of Pose and Site:

$$Pose_1 = \text{Perpendicular}; Pose_2 = \text{Oblique}$$

and

$$Site_j \in \{Vehicle, Barricade_1, \dots, Barricade_3, Logs\}$$

$$Distance_{i,j} = \mu + Site_j + Pose_i + \epsilon_{i,j}$$

where epsilon is error from unmodeled effects. In Table 2, we include results from an analysis of all runs, using a value of 0 for the Max Correct Detection distance when the barrier was not correctly detected at all. The results may be interpreted as the Site selected may not be significant in determining the detection distance, but given a particular site, the Pose was. The second analysis (Table 3) includes only runs in which the barrier was detected, omitting those in which detection never occurred. The results may be interpreted as the Site selected may not be significant in determining the detection distance, and neither may be the Pose. This would support possibility (A), since without the failures you cannot find a difference based on angle of approach. One caveat is that we did not set out to test this as a hypothesis and had we done so, we would have focused on angles and distances at which the barrier was always in view during the run. A second caveat is that with fewer observations, an effect is harder to detect, so using 13 runs instead of 19 would, by itself, make finding an effect from Pose less likely. A third analysis is shown in Table 4.

$$Pose_1 = \text{Perpendicular}; Pose_2 = \text{Oblique}$$

$$Success \in \{Yes, No\}$$

and

$$Site_j \in \{Vehicle, Barricade_1, \dots, Barricade_3, Logs\}$$

$$Distance_{i,j} = \mu + Success + Site_j + Pose_i + \epsilon_{i,j}$$

In this model, success is included as a variable, which should be significant because failed means the maximum distance is 0. Including this variable, though, allows us to pose the question this way: Given success, is the Site significant, and given success and the site, is the Pose significant? The answer is that given a successful detection, we do not find evidence that Site, or Site and Pose, matter.

Table 2 Analysis of variance for 19 runs, including nondetections

Factor	P value	Description
Pose	<0.02	Perpendicular or oblique
Site	>0.10	Barrier type: vehicle, barrier 1, 2, 3, or logs

Table 3 Analysis of variance for 13 runs, omitting nondetections

Factor	P value	Description
Pose	>0.2	Perpendicular or oblique
Site	>0.7	Barrier type: vehicle, barrier 1, 2, 3, or logs

Table 4 Analysis of variance for 19 runs, including nondetections

Factor	P value	Description
Success	<0.01	Expected since “No” means the distance must be zero
Pose	~0.2	Perpendicular or oblique
Site	>0.9	Barrier type: vehicle, barrier 1, 2, 3, or logs

The results are summarized as follows:

- The barrier detection algorithm functioned on the platform.
- Detection generally (except in 1 run) did not occur at greater than 45 m.
- Detections that researchers regarded as correct occurred within 15 m.
- The algorithm was more likely to fail to detect the barrier from an oblique angle, but this may be a limitation of the combination of limited camera field of view and algorithm together, rather than a failure of just the algorithm.
- In cases in which a correct detection occurred, we did not find evidence that the angle of approach affected the distance at which the detection occurred.
- We did not find evidence that the detection distance varied because of the type of barrier, although we limited our barriers to those that we knew would be detected as objects.

5. Conclusion

The experiment took place over the course of one day, with 19 runs being analyzed. This analysis suggests that the barrier detection algorithm generally correctly detects barriers at ranges of not more than 15 m in the experimentation area, consistently across type of barrier. There may be a difference based on angle of approach to the barrier, but it probably is because of a failure to detect the barrier. That difference could be because at the present detection range of the algorithm, the barrier is out of field of view when approached at an oblique angle.

6. References

- Deep Green. Army data science competition. Army Office of Business Transformation; n.d. [accessed 2023 July 18]. <https://www.obtportal.army.mil/EDA/DeepGreen/>.
- Ellis ST, Harrison AV. Detecting functional objects using multi-modal data. Proc. SPIE 12538, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications V; SPIE Defense + Commercial Sensing; 2023 Apr–May 5; Orlando, FL. doi:10.1117/12.2664059.
- Google Maps. R2C2 experimentation facility [Map]. Google Maps; n.d. [accessed 2023 Feb 8]. <https://www.google.com/maps/@39.3509434,-76.3448752,138m/data=!3m1!1e3>.
- Hodges J. Functional and physical object characteristics and object recognition in improvisation. *Comput Vis Image Underst.* 1995;62(2);147–163.
- Jiang P, Osteen P, Wigness M, Saripalli S. RELLIS-3D dataset: data, benchmarks and analysis. 2021 IEEE International Conference on Robotics and Automation (ICRA); 2021 May 30–June 5; Xi'an, China. p. 1110–1116. doi:10.1109/ICRA48506.2021.9561251.

Appendix. Complete Record of Results

This appendix summarizes the data collected from the experiment described in this report. Table A-1 contains the record of all runs analyzed in the report. The values of the table are described as follows.

Run: indicates the number of the experimental run.

Distance: distance in meters from the obstacle (barrier).

Pose: refers to the direction of approach, which could be either closer to perpendicular or closer to oblique and is shown in the figures corresponding to each run.

Well-defined regions: whether the semantic segmentation algorithm divided the terrain and objects into well-defined regions, which would allow terrain to be recognized as traversable and for objects to be recognized as blocking the traversable terrain (yes) or failed to do so (no).

Regions correctly labeled: whether the semantic segmentation algorithm labeled the regions with a correct label (yes) or did not (no).

Detected barrier: whether the barrier detection algorithm detected the object as an obstacle (yes or no).

Correct calculation: whether the detection of the barrier was correct given the semantic segmentation and class labels applied during that segmentation, in the opinion of the researcher (Table A-1). This calculation was based on the labeled regions (Rgn) identified by semantic segmentation.

Table A-1 Complete list of runs and results for each run

Run	Distance (m)	Pose	Well-defined regions	Regions correctly labeled	Detected barrier	Correct calculation
1	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	YES	NO - Rgn to left not blocked
	30		YES	YES	NO	YES
	15		YES	YES	YES	YES
	10		YES	YES	YES	YES
	5		N/A	N/A	N/A	N/A
2	60	Angle	YES	NO	NO	NO
	45		YES	NO	NO	NO
	30		NO	NO	NO	NO
	15		NO	YES	YES	YES
	10		Out of range	Out of range	Out of range	Out of range
	5		Out of range	Out of range	Out of range	Out of range
3	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	YES
	15		NO - Edges of bldg. are off	NO - Bldg mapped to sky	NO	NO -Rgn Bisected
	10		NO - Edges of bldg. are off	NO - Bldg mapped to sky	YES	YES
	5		Out of range	Out of range	Out of range	Out of range
4	60	Angle	YES	NO - Trees mapped to Bldg.	NO	YES
	45		YES	NO - Trees mapped to Bldg.	YES	NO - Rgn to right not blocked
	30		YES	NO - Trees mapped to Bldg.	NO	YES
4	15		YES	YES	YES	NO - Rgn to right not blocked
	10		YES	YES	YES	NO - Rgn to right not blocked
	5		Out of range	Out of Range	Out of Range	Out of Range
5	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	YES	YES
	30		YES	YES	NO	NO - Rgn Bisected
	15		YES	YES	YES	YES
	10		YES	YES	YES	YES
	5		YES	YES	YES	YES
6	60	Angle	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	YES
	15		YES	YES	NO	YES
	10		YES	YES	NO	YES
	5		YES	YES	NO	YES
7	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	NO - Rgn Bisected
	15		YES	YES	YES	YES
	10		YES	YES	YES	YES
	5		YES	YES	YES	YES
8	60	Angle	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	YES
	15		YES	YES	NO	YES
	10		YES	YES	NO	YES

Table A-1 Complete list of runs and results for each run (continued)

Run	Distance (meters)	Pose	Well-defined regions	Regions correctly labeled	Detected barrier	Correct calculation
	5		YES	YES	YES	NO - Rgn to left not blocked
9	Begin	Contour	YES	YES	NO	YES
	Barrier 3		NO - Rgns. shape off	NO - Terrain classified to non-traversable class	YES	NO - Rgn to right not blocked
	Barrier 2		YES	YES	YES	NO - Rgn to right not blocked
	Barrier 1		YES	YES	YES	NO - Rgn to right not blocked
	NA		NA	NA	NA	NA
	NA		NA	NA	NA	NA
10	Begin	Contour	YES	YES	NO	YES
	Barrier 1		YES	YES	NO	YES
	Barrier 2		YES	YES	NO	YES
	Barrier 3		YES	YES	NO	YES
	NA		NA	NA	NA	NA
	NA		NA	NA	NA	NA
11	Begin	Contour	YES	YES	YES	YES
	Barrier 4		YES	YES	NO	NO - Rgn Bisected
	Barrier 3		YES	YES	YES	YES
	Barrier 2		YES	YES	NO	NO - Rgn Bisected
	Barrier 1		YES	YES	NO	YES
	NA		NA	NA	NA	NA
12	Begin	Contour	YES	YES	YES	YES
	Barrier 1		YES	YES	NO	YES
	Barrier 2		YES	YES	NO	NO
	Barrier 3		YES	YES	NO	YES
	Barrier 4		YES	YES	YES	YES
	NA		NA	NA	NA	NA
13	60	Angle	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	NO - Rgn Bisected
	15		YES	YES	NO	NO - Rgn Bisected
	10		YES	YES	YES	YES
	5		YES	YES	YES	YES
14	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	YES
	15		YES	YES	YES	YES
	10		Ended early	Ended early	Ended early	Ended early
	5		Ended early	Ended early	Ended early	Ended early
15	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	YES	YES
	30		YES	YES	YES	YES
	15		YES	YES	YES	YES
	10		YES	YES	YES	YES
	5		Ended early	Ended early	Ended early	Ended early
16	60	Angle	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	NO - Rgn Bisected

Table A-1 Complete list of runs and results for each run (continued)

Run	Distance (meters)	Pose	Well-defined regions	Regions correctly labeled	Detected barrier	Correct calculation
	15		YES	YES	NO	NO - Rgn Bisected
	10		YES	YES	NO	NO - Rgn Bisected
	5		YES	YES	YES	YES
17	60	Angle	YES	YES	YES	YES
	45		YES	YES	NO	YES
	30		YES	YES	YES	NO - Obstacle not barrier
	15		Out of range	Out of range	Out of range	Out of range
	10		Out of range	Out of range	Out of range	Out of range
	5		Out of range	Out of range	Out of range	Out of range
18	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	NO - Rgn Bisected
	15		YES	YES	NO	NO - Rgn Bisected
	10		YES	YES	YES	YES
	5		YES	YES	YES	YES
19	60	Angle	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	NO - Rgn Bisected
	15		YES	YES	YES	YES
19	10		NO - Bad edges	NO - Bldg mapped to sky	YES	YES
	5		Out of range	Out of range	Out of range	Out of range
20	60	Perpendicular	YES	YES	NO	YES
	45		NO - Bldg. Disappeared	NO - Bldg mapped to sky	NO	YES
	30		NO - Bldg. Disappeared	NO - Bldg mapped to sky	NO	NO - Rgn Bisected
	15		YES	NO - Misclassification	YES	YES
	10		NO - Bldg. Disappeared	NO - Bldg mapped to sky	YES	YES
	5		Ended early	Ended early	Ended early	Ended early
21	60	Perpendicular	NA - no output	NA - no output	NA - no output	NA - no output
	45		NA - no output	NA - no output	NA - no output	NA - no output
	30		NA - no output	NA - no output	NA - no output	NA - no output
	15		NA - no output	NA - no output	NA - no output	NA - no output
	10		NA - no output	NA - no output	NA - no output	NA - no output
	5		NA - no output	NA - no output	NA - no output	NA - no output
22	60	Perpendicular	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	NO	NO - Rgn Bisected
	15		YES	YES	NO	NO - Rgn Bisected
	10		Out of range	Out of range	Out of range	Out of range
	5		Out of range	Out of range	Out of range	Out of range
23	60	Angle	YES	YES	NO	
	45		YES	YES	YES	NO - Rgn not Bisected
	30		YES	YES	NO	YES
	15		YES	YES	NO	YES

Table A-1 Complete list of runs and results for each run (continued)

Run	Distance (meters)	Pose	Well-defined regions	Regions correctly labeled	Detected barrier	Correct calculation
	10		YES	YES	NO	YES
	5		NO - Log shape off	NO - Did not detect log	NO	YES
24	60	Angle	YES	YES	NO	YES
	45		YES	YES	NO	YES
	30		YES	YES	YES	NO - Rgn not Bisected
	15		YES	YES	NO	YES
	10		YES	YES	NO	YES
	5		Out of range	Out of range	Out of range	Out of range

List of Symbols, Abbreviations, and Acronyms

2-D	two-dimensional
AIMM ERP	Artificial Intelligence for Maneuver and Mobility Essential Research Program
ARL	Army Research Laboratory
Bldg.	building
DEVCOM	US Army Combat Capabilities Development Command
LiDAR	laser imaging, detection, and ranging
NA	not applicable
R2C2	Robotics Research Collaboration Campus
RGB	red, green, blue
Rgn	Region identified by the barrier detection algorithm

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

1 DEVCOM ARL
(PDF) FCDD RLB CI
TECH LIB

5 DEVCOM ARL
(PDF) FCDD RLA JC
M CHILDERS
L HERNANDEZ
C LENNON
FCDD RLA IB
S ELLIS
A HARRISON