Exploring Adversarial Obstacle Attacks in Search-based Path Planning for Autonomous Mobile Robots

Adrian Szvoren¹, Jianwei Liu¹, Dimitrios Kanoulas^{1*}, Nilufer Tuptuk^{2*}

Abstract—Path planning algorithms, such as the search-based A*, are a critical component of autonomous mobile robotics, enabling robots to navigate from a starting point to a destination efficiently and safely. We investigated the resilience of the A* algorithm in the face of potential adversarial interventions known as obstacle attacks. The adversary's goal is to delay the robot's timely arrival at its destination by introducing obstacles along its original path.

We developed malicious software to execute the attacks and conducted experiments to assess their impact, both in simulation using TurtleBot in Gazebo and in real-world deployment with the Unitree Go1 robot.

In simulation, the attacks resulted in an average delay of 36%, with the most significant delays occurring in scenarios where the robot was forced to take substantially longer alternative paths. In real-world experiments, the delays were even more pronounced, with all attacks successfully rerouting the robot and causing measurable disruptions. These results highlight that the algorithm's robustness is not solely an attribute of its design but is significantly influenced by the operational environment. For example, in constrained environments like tunnels, the delays were maximized due to the limited availability of alternative routes.

I. INTRODUCTION

With the recent focus on Artificial Intelligence (AI) across several fields, the adoption and development of robots and other AI-based cyber-physical systems have become much faster and easier, leading to widespread use of them.

While these technological advances have been significant, the focus has primarily been on functionality and innovations, leaving safety and security considerations, particularly in robotics, considerably behind. The environments where the robots are being deployed, such as logistics, public spaces and streets, and disaster areas, are rarely threat-free and can be subject to attacks from adversaries. While some studies have been conducted, the vulnerability of mobile robots to adversarial attacks remains largely underexplored. The need for cyber-security and cyber-safety in robotics differs from securing traditional computing systems due to the added dimension of physical interaction [1]. This added aspect means that gaining control over a robot could result in

This work was supported by the UKRI FLF [MR/V025333/1] (RoboHike) and the CDT in Cybersecurity [EP/S022503/1]. For the purpose of Open Access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

harm, not only to the robot itself but also to its immediate environment, posing a potential and serious risk of injury to nearby individuals.

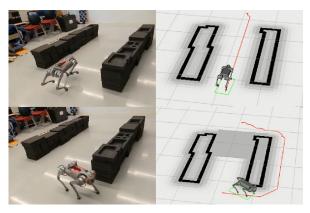


Fig. 1: Real-world deployment without any attacks (top) and with the obstacle attack (bottom), showing the planned path in red. The attack does not change the physical environment, only the robot's cost map.

Given the threats faced by autonomous mobile robots in their operational environments, it is critical to develop robust path planning algorithms that can withstand adversarial manipulations. This study examines the robustness of search-based A* path planning against adversarial manipulations in the operational environment. It explores the implications for safety and security in autonomous robotics using a methodologically simple yet effective 'obstacle attack'. While previous studies [2], [3], [4] have focused on completely immobilising or severely restricting the agent's movement, we focus on causing delays in the agent's pathfinding efforts, thereby reducing its likelihood of detection by defensive mechanisms. This is achieved by determining obstacles' optimal placement and subsequent integration into the operational environment to delay the robot's timely arrival at its destination. We conducted a set of experiments, both simulation within a static warehouse setting and on a physical robot in small environments, and collected data to demonstrate the effectiveness of the attacks. We studied the time required for the robotic agent to reach its goal and the strategic placement of obstacles.

The paper is organised as follows: Section II presents a comprehensive literature review, discussing a variety of adversarial attacks targeting robotic agents and the existing countermeasures; Section III outlines the threat model and attack used in this study, describing the experimental setup

¹Adrian Szvoren, Jianwei Liu, and Dimitrios Kanoulas are with the Department of Computer Science, University College London, Gower Street, WC1E 6BT, London, UK. {adrian.szvoren.23, jianwei.liu.21, d.kanoulas}@ucl.ac.uk

²Nilufer Tuptuk is with the Department of Security and Crime Science, University College London, Gower Street, WC1E 6BT, London, UK. {n.tuptuk}@ucl.ac.uk

^{*} Equal contribution.

for both simulation and real-world deployment, as well as the data collection techniques; Section IV details the data gathered from the experiments, providing a quantitative foundation for subsequent analysis; Section V provides an interpretation of the data with conclusions regarding the algorithm's resilience; and Section VI concludes the paper, summarising the research findings, drawing implications for future developments in autonomous robotics.

II. RELATED WORK

Path planning has been heavily studied in the past [5], [6], [7], [8]. Here, we briefly review the work related to adversarial attacks in robots.

While there have been some studies on the security threats that affect robotic hardware, firmware/OS, and applications [9], [10], [11], [12], less focus has been given to studying the vulnerabilities of AI deployed in robotics [13]. Robots interacting with a physical environment bring a new category of threats beyond software and hardware risks referred to as behaviour threats [14]. For example, an autonomous delivery robot can make nearby pedestrians feel threatened, prompting them to take action to render the robot immobile [15].

Adversarial attacks can significantly impact autonomous robots, disrupting control systems and decision-making capabilities. Vemprala et al. [2] investigate the susceptibility of iterative optimisation-based trajectory planners to adversarial attacks, showing that such attacks can disrupt the optimisation process by manipulating the cost function and lead to failures or increased computation times. The authors test their method by applying adversarial strategies to two state-of-the-art trajectory planners, demonstrating that these planners are consistently vulnerable to such attacks. Cavorsi et al. [3] present an algorithm for resilient path planning in robot teams, where robots can coordinate despite adversaries. The authors use Control Barrier Functions (CBFs) to balance resilience and safe navigation, developing a method to identify and handle trade-offs between these objectives. Wu et al. [4] focus on the vulnerability of robot motion planning algorithms to adversarial attacks, specifically in the physical environment. Their study highlights two types of attacks: planner failure and blindspot attacks. Planner failure attacks involve making subtle changes to the physical environment, causing the motion planner to fail to find a valid path. In contrast, blindspot attacks exploit occlusions and sensor field-of-view limitations, leading the planner to generate a collision-free trajectory that intersects with unperceived obstacles. The reported experimental results indicate that these attacks can drastically increase failure rates to 95%and collision rates to 90% with only minor environmental modifications.

Clark et al. [16] examine the manipulation of an autonomous robotic vehicle's ultrasonic collision avoidance sensors to redirect its path. By employing the Q-learning algorithm for real-time route control, they illustrate the potential for indirect attacks on robotic systems, emphasising the need for robust countermeasures. Shi et al. [17] explore

the vulnerabilities of learning-based locomotion controllers in quadrupedal robots using adversarial attacks. Their study focuses on the robustness of controllers that employ deep Reinforcement Learning (RL) against real-world uncertainties like sensor noise and external perturbations. Despite empirical robustness, this research shows that these controllers can fail significantly under carefully crafted, low-magnitude adversarial sequences.

Lechner et al. [18] studied the robustness-accuracy tradeoff, exploring the balance between maintaining accuracy and ensuring robustness against adversarial attacks. Focusing on two neural controllers - one for vision-based commands and another for LiDAR-based motor control; their findings suggest that adversarial training can introduce new errors, potentially undermining the robot's learned behaviour. Lechner et al. [19] revisit the robustness-accuracy trade-off, incorporating advancements like overparametrisation and vision transformers. Their study shows that integrating adversarial weight perturbation reduces the gap between robustness and accuracy.

To the best of our knowledge, no previous studies have examined the robustness of path planning algorithms under adversarial manipulations in operational environments where obstacles are introduced along the intended path to delay timely arrival at the destination.

III. METHODOLOGY

In this section, the threat model will first be introduced, covering the adversary's objectives, incentives, knowledge, and capabilities. It then introduces the obstacle attack and experimental setup to show how the attacks were implemented in simulation and validated real-world deployments, along with the data collection process.

A. Threat Model

Objectives: The adversary's objective is to delay the physical agent's timely arrival at its destination by introducing obstacles along its original path, whether physical or virtual, thereby forcing it to take less optimal alternative routes. If necessary, the adversary could modify the attack to direct the agent along a specific path or through a particular point on the map, transforming the initial non-targeted attack into a targeted one. The adversary's goal is to ensure that the attack remains as undetectable as possible, limiting it to adding a single obstacle.

Incentives: The adversary's motivations predominantly hinge on inflicting financial harm. This includes tactics such as delaying or diverting delivery robots for potential theft opportunities and inflicting reputational damage on companies that operate robotic systems or directly on the robotic systems themselves. Furthermore, these malicious actions could extend to risking human lives, especially concerning robots employed in critical roles like disaster relief and rescue services. The extent of these actions not only threatens the operational integrity and trust in robotic solutions but also emphasises the need for robust security measures and ethical

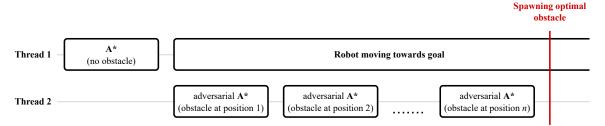


Fig. 2: Parallel execution of the obstacle attack.

guidelines to mitigate such risks and safeguard both property and lives.

Knowledge & Capability: We assume that the adversary could carry out the attack virtually by manipulating the agent's cost map to create a false obstacle. This manipulation could be achieved through a man-in-the-middle attack [20] or malicious software. To calculate the optimal obstacle placement, the adversary would need access to the map of the environment in which the agent operates. We assume that the adversary possesses this knowledge, as well as the agent's starting (current) and goal positions. These assumptions are reasonable, given that robots repeatedly operate in the same environment.

B. Obstacle Attack

The attack described employs a method we named bruteforce obstacle attack. This approach methodically tests every feasible location for placing an obstacle along the path initially planned. The process is summarised by pseudo-code in Algorithm 1.

Algorithm 1 Obstacle Attack

```
path ← A*(start, goal, map)
path<sub>longest</sub> ← path
for step in path do
path<sub>new</sub> ← run A* with obstacle at step
if LENGTH(path<sub>new</sub>) > LENGTH(path<sub>longest</sub>) then
obstacle<sub>best</sub> ← obstacle at step
path<sub>longest</sub> ← path<sub>new</sub>
end if
end for
add obstacle<sub>best</sub> to map
```

The process begins by determining the actual path that the agent should follow. This is achieved by deploying the global path planning algorithm, in our case A*, which requires the starting point, destination (goal), and cost map as inputs. Following this initial use of the A* algorithm to establish the baseline path, the algorithm is executed repeatedly, equivalent to the number of steps identified within the original path. During these subsequent iterations, an obstacle is introduced into the cost map at different positions. This technique makes it feasible to ascertain the most strategic placement of the obstacle that yields the maximum path extension required for the agent to reach its goal.

The size of the obstacle is carefully chosen to balance two objectives: the obstacle must be large enough to force the agent to take a significantly longer alternative route, but it should not be so large as to completely block the agent's path, which could render the scenario unrealistic or trivial.

The location of the obstacle is determined through a bruteforce search across all possible positions along the agent's path. For each position, the algorithm evaluates the resulting path extension and selects the placement that maximizes the delay. This iterative process ensures that the obstacle is positioned in the most disruptive yet plausible location, aligning with the adversary's goal of maximizing impact while maintaining undetectability.

Once the optimal position for the obstacle has been identified through this iterative process, the obstacle is inserted in the agent's cost map, thus impacting the path's overall length and complexity. The effectiveness of this method lies in its simplicity and the brute-force nature of testing every possible variation to find the most obstructive placement of the obstacle.

For the A* algorithm, the larger the map, the longer it takes to find a feasible path. Running the obstacle attack algorithm sequentially would, therefore, cause a delay between receiving the goal position and commencing the motion towards it. To address this, the attack algorithm has been designed to run in parallel with the normal function of the Robot Operating System (ROS) [21] node. Fig. 2 illustrates the two threads utilised in the attack process. Thread 1 encompasses the intended sequence of execution, wherein it first determines the path to the goal and then initiates the robot's movement along the specified path towards the goal. Thread 2 is commenced after the path planning algorithm is completed. It is where the adversarial A* algorithm runs once for each position along the path, and upon evaluating the last position, it deploys the obstacle in the optimal position. As thread 2 runs concurrently with the robot's movement, the robot may have already traversed the obstacle position; this is one type of data we evaluate in the experiments.

C. Experimental Setup

In the following sections, we explain how the attacks were created first through simulation using path planning simulated in Gazebo and later validated using real-world deployment on a Unitree Go1 robot ¹.

¹https://www.unitree.com/go1/

1) Simulation: The experiments were conducted in Gazebo simulation environment using the ROS. An existing project has been selected as a reference for path planning simulated in Gazebo. The project features a ROS implementation of various path planning algorithms, a simulated warehouse environment, and a robotic agent. Specifically, we used the TurtleBot 3 Waffle as our robotic agent, equipped with the 360 Laser Distance Sensor LDS-01, a 2D laser scanner capable of 360-degree sensing. This sensor collects data around the robot for use in SLAM. Our global path planning algorithm was A*, while we employed a PID controller for local path planning. To carry out the attack, we deployed a malicious ROS node as the attack vector.

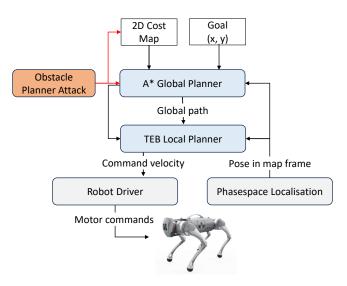


Fig. 3: Navigation stack for the real robot deployment.

2) Real-world Deployment: A schematic representation of the real-world deployment of the obstacle attack is depicted in Fig. 3. We validate the performance of the attack in real-world conditions on a Unitree Go1. To leverage the existing robot navigation frameworks, the attack is integrated with the 2D ROS navigation stack to act as a part of the global path planner. Given the cost map, the attack generates an obstacle in the global path, making sure the obstacle does not fully block all the possible paths. The local planner and obstacle avoidance algorithm used is the Timed-Elastic-Band (TEB) [22], [23] with localisation provided by the Phasespace tracking system in real-time at 960 Hz. Examples of successful deployment of the pipeline can be seen in Fig. 1.

D. Data Collection

To collect data, we conducted the experiments in simulation and real-world deployment. **Simulation scenario**: We positioned the robot in the centre of the warehouse and instructed it to travel to one of the 23 specified locations, where the focus was on end-to-end navigation in larger environments. **Real-world scenarios**: We conducted three scenarios with a real robot, each scenario with a different

map (tunnel, turn, and t-wall), testing the robot's performance in simpler, smaller environments.

Both the simulation and real-world experiments were repeated three times for each map/start/end without any interference (known as benign conditions) and three times for each map/start/end with interference (known as adversarial conditions), resulting in 157 experimental runs. The data we collected included the coordinates of the start and end locations, the time taken by the robot to reach them, and for runs with interference, the time taken to introduce the obstacle, the coordinates of the obstacle, and whether the obstacle was placed before the robot reached the obstacle's location.

IV. RESULTS

This section presents the results from the data collection for both simulated and real-world experiments.

Fig. 4 compares simulation scenarios with predetermined goals and dynamically placed obstacles, where it can be seen that the attack places obstacles in narrow pathways and on the corners to delay the robot the most.

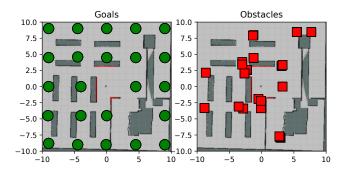


Fig. 4: Goal positions used in the experiments (left) and all the calculated obstacle positions (right).

Fig. 5 presents the simulation results comparing the performance of a robotic system under benign and adversarial conditions. The central scatter plot illustrates the relationship between the Euclidean distance to a goal and the time taken to reach that goal, capturing all the data points from the experiment. In the benign condition, the robot navigates through the environment with no interference from the adversary. The data points exhibit a clear upward trend, where the time to reach the goal increases with the Euclidean distance. In contrast, in the adversarial condition, the robot encounters an obstacle that complicates its path, affecting its efficiency. The adversarial trend is less steep than the benign trend, suggesting that although obstacles introduce delays, the robot can somewhat mitigate the impact on travel time across varying distances, opening the possibility of adaptive path finding strategies that limit the increase in time as distance grows.

A critical factor influencing the success of the attack is the timing of obstacle injection relative to the agent's progress along its path. Obstacles injected early in the path give the agent more opportunities to replan efficiently, often resulting

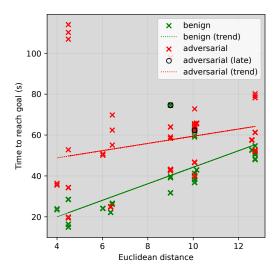


Fig. 5: A comparison of robotic navigation times under benign and adversarial conditions, showing the impact of obstacles on travel time relative to Euclidean distance.

in smaller delays. For example, if an obstacle is placed near the start of the path, the agent can explore alternative routes with minimal backtracking, reducing the overall impact of the attack.

Fig. 6 consists of two plots that analyse the timing and success of obstacle spawning in the simulated experiments. The high success rate of 91.3% indicates that the adversarial A* algorithm is generally effective at identifying optimal injection times to maximize delays. However, the 8.7% failure rate highlights instances where obstacles were injected too late or in suboptimal locations, allowing the agent to bypass them with minimal disruption. These failures underscore the importance of precise timing in adversarial attacks and suggest that further refinement of the injection strategy could improve the algorithm's effectiveness.

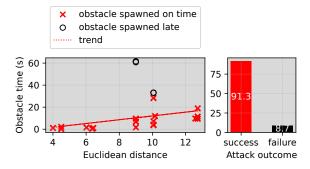


Fig. 6: The relationship between Euclidean distance, obstacle spawn time, and the success rate of obstacle attacks.

Fig. 7 presents a detailed comparison of the performance of the A* path planning algorithm under two distinct conditions: "benign" navigation and "adversarial" navigation, where obstacles were intentionally introduced as part of an obstacle attack strategy. The figure captures both the mean absolute (top subfigure) and percentage differences (bottom

subfigure) in navigation time between these two scenarios, providing insights into the effectiveness of the obstacle attacks. The mean time delay across all goal positions is 17.78 seconds, while the percentage delay across all experiments is 36.01%.

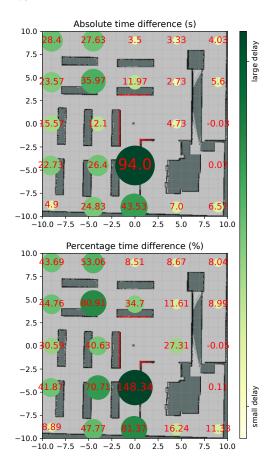


Fig. 7: Comparison of mean absolute and percentage difference in time taken between benign and adversarial scenarios for each goal position.

The real-world experiments show performance on smallerscale maps compared to the simulation. This was to investigate the effects of the obstacle attack on robots operating in environments with limited mobility. The maps used in these experiments can be seen in Fig. 8, and it shows both the benign and adversarial scenarios.

The results are summarized in Table I. The attack caused the largest delay in the tunnel scenario, where the obstacle completely redirected the robot, resulting in an average delay of 10 seconds. This scenario represents the worst-case performance for the agent, as the constrained environment limited the robot's ability to find efficient alternative paths. In the turn scenario, the delays were minimal, with the robot's travel time being almost the same as without the attack. This suggests that the impact of adversarial obstacles depends heavily on the environment's layout and the availability of alternative routes. In the final scenario, where the robot was not fully redirected due to the lack of alternative paths, the attack still managed to delay the robot's arrival time by almost

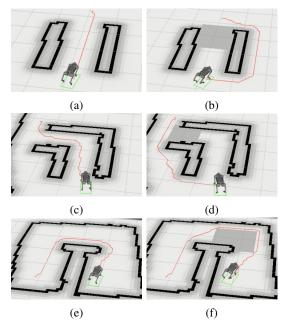


Fig. 8: The 3 maps used in real-world experiments: tunnel (a)(b), turn (c)(d), and t-wall (e)(f). The first column shows the original path taken without an attack, while the second column shows the calculated obstacle placement (grey square) and the alternative path taken by the robot.

10%. This demonstrates that even in less ideal conditions for the adversary, the A* algorithm remains susceptible to delays caused by strategically placed obstacles.

Scenario	Map	Mean absolute delay	Mean percentage delay	Success rate
simulation	warehouse	17.78s	36.01%	91.3%
real- world	tunnel	10.07s	86.08%	100%
	turn	0.23s	1.75%	100%
	t-wall	1.57s	9.65%	100%
	overall	3.95s	32.49%	100%

TABLE I: Experimental results.

V. DISCUSSION

The experiments revealed that adversarial obstacles can cause significant delays, with the worst-case delays reaching up to 10 seconds in the tunnel scenario. These delays highlight the vulnerabilities of the current search-based A^* path planning algorithm to adversarial conditions. While the success rate of the attacks was high (91.3% in simulation and 100% in real-world scenarios), the magnitude of the delays provides a more meaningful measure of the algorithm's susceptibility. For example, in environments with limited alternative routes, such as tunnels, the delays were maximized, underscoring the algorithm's dependence on environmental factors. These findings emphasize the urgent need for more robust path planning algorithms capable of withstanding adversarial attacks.

From the trends associated with mean absolute and percentage delays, we observed that traversal times increase with distance under benign and adversarial conditions. However, it is noteworthy that the trend line is less steep in adversarial conditions, signalling a relatively uniform impact of obstacles irrespective of distance. This observation led us to speculate on the nature of the challenges faced by the robot in navigational tasks. Despite the introduction of obstacles presenting a consistent difficulty level, our findings suggest that the robot's performance does not significantly deteriorate over longer distances when obstacles are present. This could imply the existence of a "threshold effect", where beyond a certain level of introduced complexity, further increments in distance do not correspond with a disproportionate increase in navigational difficulty.

Another important observation was the high impact of the attack strategies in constrained spaces, specifically on the left side of the simulated map. This area, having limited space for motion, saw significant delays. The restricted manoeuvrability in such areas severely compromises the robot's ability to adapt to and circumvent obstacles, leading to notable increases in delay times. By carefully considering the spatial characteristics of the environment, one can maximise the impact of adversarial conditions, pointing to a nuanced understanding of how space and adversarial strategies interact.

VI. CONCLUSION

We investigated the resilience of the search-based path planning algorithm A* in the face of potential adversarial interventions known as obstacle attacks. We introduced the threat model, defined obstacle attack as a brute-force effort and tested its efficacy through both simulation and real-world deployment. The findings shows that the algorithm's robustness is not solely an attribute of its design but is significantly influenced by the operational environment. Specifically, factors such as the size of the environment and the proximity to the designated goal played crucial roles.

While this study focuses on adding a single obstacle to simplify the analysis, the approach is justified by the need to establish a baseline understanding of the adversarial impact. A single obstacle provides a clear and measurable way to evaluate the worst-case delay without introducing additional variables, such as varying obstacle sizes or quantities. However, we acknowledge that in real-world scenarios, an adversary might employ multiple smaller obstacles to achieve a similar delay while maintaining a lower profile. For example, several small obstacles could be strategically placed to create a cumulative delay without being as visually conspicuous as a single large obstacle.

This research contributes to a better understanding of the vulnerabilities inherent in the widely used A* path planning algorithm. It demonstrates the importance of considering environmental variables in assessing algorithmic robustness. Furthermore, the successful execution of the obstacle attack in many test scenarios highlights a critical area for future research: improving the security of robotic systems by developing more resilient algorithms to protect autonomous systems against such adversarial manipulations.

REFERENCES

- D. Portugal, S. Pereira, and M. S. Couceiro, "The role of security in human-robot shared environments: A case study in ros-based surveillance robots," in 2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). IEEE, aug 2017
- [2] S. Vemprala and A. Kapoor, "Adversarial attacks on optimization based planners," in 2021 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2021, pp. 9943–9949.
- [3] M. Cavorsi, L. Sabattini, and S. Gil, "Multi-robot adversarial resilience using control barrier functions," *IEEE Transactions on Robotics*, 2023.
- [4] W. Wu, L. Chen, Y. Yuan, and W. Zhao, "An adaptive hierarchical path planning algorithm for autonomous robots in unknown dynamic environments," in 2024 IEEE International Conference on Robotics and Automation (ICRA). IEEE, may 2024.
- [5] K. Ellis et al., "Navigation among movable obstacles with object localization using photorealistic simulation," in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022.
- [6] J. Liu, S. Lyu, D. Hadjivelichkov, V. Modugno, and D. Kanoulas, "ViT-A*: Legged Robot Path Planning using Vision Transformer A*," in 2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids). IEEE, 2023, pp. 1–6.
- [7] J. Liu, M. Stamatopoulou, and D. Kanoulas, "DiPPeR: Diffusion-based 2D Path Planner applied on Legged Robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [8] M. Stamatopoulou, J. Liu, and D. Kanoulas, "DiPPeST: Diffusion-based Path Planner for Synthesizing Trajectories Applied on Quadruped Robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [9] G. W. Clark, M. V. Doran, and T. R. Andel, "Cybersecurity issues in robotics," in 2017 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA). IEEE, mar 2017.
- [10] O. M. Creado and P. D. Le, "Enforcing security in artificially intelligent robots using monitors (short paper)," in *Lecture Notes in Computer Science*. Springer International Publishing, 2017, pp. 648–659.
- [11] K. A. Yousef, A. AlMajali, S. Ghalyon, W. Dweik, and B. Mohd, "Analyzing cyber-physical threats on robotic platforms," *Sensors*, vol. 18, no. 5, p. 1643, may 2018.
- [12] J.-P. A. Yaacoub, H. N. Noura, O. Salman, and A. Chehab, "Robotics cyber security: vulnerabilities, attacks, countermeasures, and recommendations," *International Journal of Information Security*, vol. 21, no. 1, pp. 115–158, mar 2021.
- [13] R. Calo, I. Evtimov, E. Fernandes, T. Kohno, and D. O'Hair, "Is tricking a robot hacking?" SSRN Electronic Journal, 2018.
- [14] M. Colledanchise, "Address behaviour vulnerabilities in the next generation of autonomous robots," *Nature Machine Intelligence*, vol. 3, no. 11, pp. 927–928, nov 2021.
- [15] J. A. Oravec, "Robo-rage against the machine: Abuse, sabotage, and bullying of robots and autonomous vehicles," in *Social and Cultural Studies of Robots and AI*. Springer International Publishing, 2022, pp. 205–244
- [16] G. Clark, M. Doran, and W. Glisson, "A malicious attack on the machine learning policy of a robotic system," in 2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE). IEEE, aug 2018.
- [17] F. Shi, C. Zhang, T. Miki, J. Lee, M. Hutter, and S. Coros, "Rethinking robustness assessment: Adversarial attacks on learning-based quadrupedal locomotion controllers," arXiv, 2024.
- [18] M. Lechner, R. Hasani, R. Grosu, D. Rus, and T. A. Henzinger, "Adversarial training is not ready for robot learning," in 2021 IEEE International Conference on Robotics and Automation (ICRA). IEEE, may 2021.
- [19] M. Lechner, A. Amini, D. Rus, and T. A. Henzinger, "Revisiting the adversarial robustness-accuracy tradeoff in robot learning," *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1595–1602, mar 2023.
- [20] M. Conti, N. Dragoni, and V. Lesyk, "A survey of man in the middle attacks," *IEEE communications surveys & tutorials*, vol. 18, no. 3, pp. 2027–2051, 2016.
- [21] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.

- [22] C. Rösmann, W. Feiten, T. Wösch, F. Hoffmann, and T. Bertram, "Trajectory Modification Considering Dynamic Constraints of Autonomous Robots," in *German Conference on Robotics (ROBOTIK)*. VDE, 2012, pp. 1–6.
- [23] C. Rösmann, W. Feiten, T. Wösch, F. Hoffmann, and T. Bertram, "Efficient Trajectory Optimization Using a Sparse Model," in *European Conference on Mobile Robots*, 2013, pp. 138–143.