USENIX

THE ADVANCED COMPUTING
SYSTEMS ASSOCIATION

WIP: Evaluation of Threats and Impacts of HD Map Tampering Attacks in Autonomous Driving

Miyu Sato and Ryunosuke Kobayashi, *Waseda University;* Kazuki Nomoto and Yuna Tanaka, *Waseda University, Deloitte Tohmatsu Cyber LLC;* Go Tsuruoka, *Waseda University;* Tatsuya Mori, *Waseda University, NICT, RIKEN AIP*

https://www.usenix.org/conference/vehiclesec25/presentation/sato

This paper is included in the Proceedings of the
3rd USENIX Symposium on Vehicle Security and Privacy.

August 11–12, 2025 • Seattle, WA, USA

978-1-939133-49-6

Open access to the Proceedings of the 3rd USENIX Symposium
on Vehicle Security and Privacy is sponsored by USENIX.

WIP: Evaluation of Threats and Impacts of HD Map Tampering Attacks in Autonomous Driving

Miyu Sato
Waseda University

Ryunosuke Kobayashi
Waseda University

Kazuki Nomoto
Waseda University,
Deloitte Tohmatsu Cyber LLC

Yuna Tanaka
Waseda University
Deloitte Tohmatsu Cyber LLC

Go Tsuruoka
Waseda University

Tatsuya Mori
Waseda University, NICT, RIKEN AIP

Abstract

High-definition (HD) maps are essential for autonomous vehicle (AV) navigation, providing detailed road and lane structure information. However, their static nature makes them vulnerable to tampering, posing significant security risks. This study systematically categorizes HD map tampering threats and evaluates their impact through an end-to-end autonomous driving simulation using Autoware and AWSIM. By modifying lane widths in HD maps, we demonstrate how small modifications can cause AVs to deviate from safe trajectories, affecting both planning and control. Our findings demonstrate the need for robust HD map verification, cryptographic validation of map updates, and a balance between HD map reliance and real-time perception. The study demonstrates the importance of securing HD maps to ensure safe and reliable AV operations.

1 Introduction

High-definition (HD) maps serve as a critical component in autonomous vehicle (AV) systems, providing precise localization and structured environmental information beyond what on-board sensors alone can achieve [8]. Unlike real-time sensor data, HD maps provide pre-processed, high-fidelity 3D representations of road layouts, lane structures, and traffic regulations to ensure robust navigation and decision making. These 3D maps include elevation changes, road slopes, and detailed structural elements such as bridges and tunnels, enhancing the AV's ability to interpret complex driving environments. Figure 1 illustrates an example of an HD map, showing a highly detailed digital representation of an intersection that helps AVs interpret road structures and navigate safely. Companies like Waymo rely on HD maps as a prerequisite for safe autonomous operation, creating detailed maps before deploying AVs in new areas [4].

However, the static nature of HD maps presents a fundamental challenge: they require continuous updates to reflect real-world changes such as road modifications, new traffic signs, or lane shifts [6]. Failure to maintain map accuracy, whether due to delayed updates or intentional tampering, can have serious consequences for AV safety.

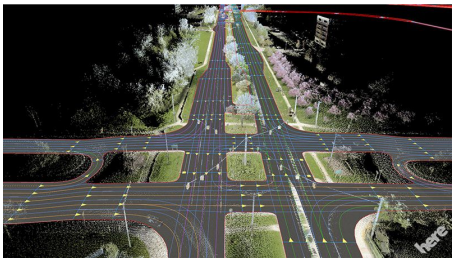

Figure 1: Example of an HD map from HERE [8], showing detailed road and lane structures for autonomous navigation.

Unlike adversarial attacks on real-time perception systems, which primarily affect individual vehicles [9, 23], compromised HD maps pose a systemic risk. Because AVs often retrieve map data from centralized servers, an attacker who modifies HD maps at the source can simultaneously affect all vehicles that rely on that data. This distinguishes HD map security threats from sensor-specific adversarial attacks, such as LiDAR spoofing [22] or adversarial patches targeting Visual SLAM [10], which affect local perception rather than system-wide navigation.

The potential impact of such discrepancies is not hypothetical; real-world incidents have demonstrated the dangers of inaccurate map data. In one reported case, a navigation system guided a vehicle onto a collapsed bridge, resulting in a fatal accident [16]. A similar risk exists in autonomous driving, where HD maps play a critical role in route planning and decision making. If an AV relies on a compromised HD map containing outdated or maliciously altered information, it may miscalculate its trajectory and navigate into dangerous areas, leading to potential accidents. Despite these risks, security research on HD map integrity remains limited compared to work on sensor attacks [9, 23] or defenses using external map cross-validation [24].

To address this gap, we define *HD map tampering attacks* as deliberate modifications of HD maps intended to mislead AVs.

This study aims to address the following research questions (RQs):

- **RQ1**: What security threats do HD map tampering attacks pose to autonomous vehicles?
- **RQ2**: How do HD map tampering attacks affect autonomous vehicle navigation and planning?

To answer RQ1, we systematically analyze HD map tampering techniques and categorize the security threats they pose. We classify attacks based on the map elements targeted, the nature of the changes, and their potential impact on AV operations. This structured analysis provides a foundation for understanding the risks posed by HD map manipulation. For RQ2, we evaluate the concrete impact of HD map manipulation through simulation. Using Autoware [17] and AWSIM [25],the open-source autonomous driving stack and a high-fidelity 3D simulator for testing autonomous vehicles, we test the HD tempering map attacks and analyze their impact on vehicle navigation. Specifically, we introduce targeted lane width modifications in HD maps and evaluate the resulting deviations in planned trajectories and actual vehicle motion.

The main contributions of this study are as follows:

- We establish a systematic framework for analyzing HD map manipulation threats, classifying attack types and their impact on AV safety.
- We conduct a simulation-based evaluation of HD map tampering, demonstrating how specific modifications affect AV planning and execution.

2 Background and Related Work

2.1 Role of HD Maps in Autonomous Driving

Autonomous driving consists of five core functions: sensing, perception, localization, planning, and control. HD maps play a critical role in several stages of this process. Sensing uses LiDAR, cameras, and radar to collect environmental data. Perception processes this data using machine learning models to detect objects and predict their movements. Localization combines sensor readings with HD map data to estimate the vehicle's precise position, often using LiDAR-based point cloud matching. Planning uses HD map information for route generation and maneuvering decisions, while control executes these plans through acceleration, braking, and steering.

Planning in autonomous driving is further divided into three levels. **Mission Planning** determines the global route using static lane information from HD maps. **Behavior Planning** ensures rule-compliant decisions such as lane changes and intersection handling. **Motion Planning** refines trajectories by optimizing for safety, comfort, and obstacle avoidance. In each of these stages, HD maps provide critical data for accurate and reliable decision-making.

2.2 HD Map Structure and Formats

HD maps consist of multiple layers that provide structured road information. Elghazaly et al. [13] categorize these into six layers:

(1) **Base Map Layer**: Represents fundamental geometric road features.
(2) **Geometric Map Layer**: Specifies detailed road geometry, including lane markings and curbs.
(3) **Semantic Map Layer**: Encodes lane attributes such as direction, speed limits, and road types.
(4) **Road Connectivity Layer**: Defines lane transitions and road network connections.
(5) **Priors Map Layer**: Contains historical traffic patterns and other learned behaviors.
(6) **Real-Time Map Data**: Provides dynamic updates on traffic conditions, weather, and road closures.

Together, these layers enable precise localization and informed planning.

The format of HD maps varies from provider to provider. OpenStreetMap (OSM) [3] serves as a widely adopted standard, with extensions such as Lanelet2 [15] tailored for autonomous driving. The Autoware framework [17] relies on Lanelet2 to represent vector maps, which were used in this study for HD map modifications.

2.3 HD Map Generation and Maintenance

The creation of HD maps requires highly accurate data collection. Mobile mapping systems (MMS) equipped with LiDAR, cameras, and GNSS sensors collect road features that are then manually processed into map data. Due to the labor-intensive nature of this process, HD map production remains costly and time-consuming. To ensure map accuracy, continuous updates are required [13]. Change detection algorithms compare new sensor data from autonomous vehicles and traffic infrastructure with existing maps. Identified discrepancies trigger updates to maintain consistency. While real-time updates are proposed [6], they require a reliable communication infrastructure, which poses deployment challenges.

2.4 Related Work

Research on HD maps includes aspects such as construction, security, and updates. Elghazaly et al. [13] and Asrat et al. [6] provide comprehensive reviews of HD map development and maintenance. Luo et al. [19] analyze security risks in localization and navigation, highlighting HD maps as a potential attack vector. Sato et al. [24] propose a defense mechanism against off-road attacks by cross-checking detected routes with map data.

Despite these contributions, research on HD map integrity remains limited compared to adversarial attacks on perception systems [9, 23]. This study aims to address this gap by

systematically analyzing HD map manipulation threats and evaluating their impact through simulation.

3 Systematic Analysis of HD Map Tampering Attacks

3.1 Threat Model

For autonomous vehicles to operate safely, HD maps must accurately reflect real-world conditions. However, they are vulnerable to both intentional tampering and outdated data, leading to navigation errors and safety risks. Tampering with HD maps can result in manipulated lane structures, falsified road connectivity, or altered traffic rules, misleading autonomous vehicles into unsafe maneuvers. Attackers may introduce incorrect speed limits, modify stop sign locations, or create nonexistent roads, potentially causing traffic disruptions or accidents. Additionally, adversaries could exploit real-time map updates by injecting false congestion data, diverting vehicles into inefficient or dangerous routes. We assume that attackers have a detailed understanding of the HD map specifications and structure, and can pre-identify road segments that target vehicles are likely to traverse. Direct access to the vehicle's internal ECUs or communication buses is not required; rather, tampering is assumed to be possible through external interfaces such as the HD map supply chain or online update pathways. The specific threat models for each attack scenario are detailed in Section 3.2.

3.2 Tampering Attack Scenarios

We present various HD map tampering attack scenarios:
Supply Chain Attacks. If a trusted source within the supply chain is compromised, falsified map data can spread to many autonomous vehicles through shared HD map updates. As HD maps expand, their management is expected to involve multiple stakeholders, including government agencies and private companies [5]. A compromised entity within this supply chain could introduce falsified map data. Once a trusted source within the supply chain is compromised, falsified map data can propagate across various systems that rely on shared HD map updates. In particular, cloud servers used for map distribution, third-party mapping services, or even authorized update nodes may serve as vectors for spreading tampered information. Given the increasing reliance on real-time map updates for safe navigation, attacks on the supply chain could have widespread and persistent effects, compromising the integrity of autonomous driving systems across broad regions.
Exploiting Update Mechanisms. Attackers can compromise HD map accuracy by exploiting the map update mechanisms themselves, both digitally and physically. HD maps are updated by aggregating sensor data from multiple autonomous vehicles, a process known as map fusion [20]. Attackers can inject false data into this process by introducing fake road signs or modifying existing ones, causing incorrect information to be incorporated into HD maps. In addition to digital tampering, attackers can physically alter the environment to manipulate the sensor input of vehicles collecting map data. For example, drawing fake lane markings on the road or placing adversarial patches can deceive perception systems, leading vehicles to misinterpret the road layout. While Sato et al. [23]demonstrated such attacks on perception systems, we suggest that similar methods could also result in incorrect data being integrated into HD maps during the update process. These types of physical-world attacks do not require insider access and exploit known vulnerabilities in camera and sensor systems.
Manipulating Open-Source Maps. Open-source HD maps allow external users to modify map data, making them vulnerable to malicious tampering by attackers. Some HD maps rely on publicly available datasets such as OpenStreetMap (OSM) [6, 7]. While these sources reduce costs, they allow anyone to edit map data. Attackers can modify road attributes or traffic signs, misleading autonomous vehicles. For example, OpenStreetMap (OSM), currently the most widely used open-source map, has been adopted in various commercial applications. Uber, for instance, has previously edited OSM data to improve routing in under-mapped areas. As autonomous driving systems increasingly adopt such open-source maps for commercial use, it is expected that more companies will rely on and potentially contribute to these platforms. Consequently, systems using editable public maps may face heightened risks of external tampering in the future, especially as their reliance on these platforms grows with broader commercial deployment. This suggests a potential vulnerability in autonomous driving systems that depend on publicly editable mapping infrastructure.

Table 1: Update Frequency of HD Map Layers.

Layer	Update Frequency
Real Time Map	Very High (Real-time)
Priors Map	High (Every few hours)
Road Connectivity	Medium (Days to weeks)
Semantic Map	Medium (Days to weeks)
Geometric Map	Medium (Days to weeks)
Base Map	Low (Weeks to months)

3.3 Attack Persistence and Update Frequency

The impact of HD map tampering depends on the update frequency of each map layer. Table 1 summarizes update intervals based on data from the Automotive Edge Computing Consortium [12]. Tampering in layers with infrequent updates, such as the Base Map, can persist for extended periods, while modifications in frequently updated layers, such as the Real-Time Map, can spread quickly and cause immediate disruptions.

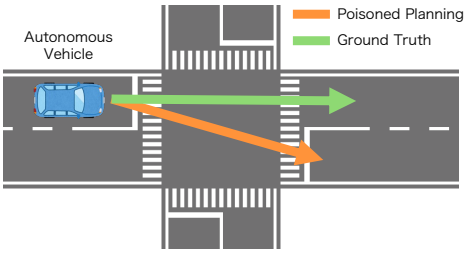

Figure 2: Illustration of Tampering Attack.

3.4 Classification of the Tampering Attacks

We systematically classify HD map tampering attacks based on affected layers, potential modifications, and their impact on autonomous vehicle modules. Table 2 categorizes these scenarios and highlights how different tampering methods can compromise the safety of autonomous driving. The impact of HD map manipulation varies depending on the feature targeted and its role in vehicle navigation and decision making. Lane and centerline modifications directly affect mission planner and localization, leading to serious navigation errors such as unintended lane changes or wrong turns. Manipulation of road attributes, such as speed limit changes or lane direction changes, can disrupt behavioral planner, potentially causing vehicles to accelerate or decelerate inappropriately or violate traffic laws.

Traffic sign modifications, including falsified locations or content, can interfere with perception-based decision making, leading to incorrect responses at intersections, stop signs, or restricted zones. In addition, the injection of fake real-time traffic data can mislead vehicles into inefficient, congested, or even dangerous routes, compromising overall traffic flow and safety. Because different HD map layers have different update frequencies and levels of reliability, the persistence and immediacy of these attacks can vary. While some modifications, such as point cloud manipulation, may have a long-lasting effect due to infrequent updates, others, such as real-time traffic data spoofing, can quickly affect multiple vehicles simultaneously. Understanding these differences is critical to assessing risk and developing effective countermeasures.

4 Evaluation of HD Map Tampering Attacks

This section presents a proof-of-concept evaluation of HD map tampering attacks using an autonomous driving simulation. Instead of conducting a comprehensive security analysis on all the possible attacks, we focus on a specific attack, namely the modification of lane width information, and demonstrate its impact on AV navigation. We describe the experimental setup, attack methodology, and evaluation results.

4.1 Experimental Overview

HD map tampering can take various forms, such as incorrect lane markings, altered traffic signals, or falsified road connectivity. Among these, we investigate a simple yet effective attack: lane width modification, where an adversary expands lane boundaries in an HD map to manipulate an AV's planned trajectory. As illustrated in Figure 2, the attack aims to mislead the AV's path planning module into deviating from its expected route. The AV, relying on the altered HD map, generates a planning trajectory that could lead it into an incorrect lane or an unsafe position on the road. We validate this hypothesis by conducting an end-to-end simulation using an open-source autonomous driving stack and a realistic urban driving scenario.

4.2 Experimental Setup

Autonomous Driving Software and Simulator. We use *Autoware* [17], a widely adopted open-source autonomous driving stack, to simulate an AV's decision-making process. To recreate a realistic driving environment, we employ *AWSIM* [25], a Unity-based autonomous driving simulator that provides a high-fidelity representation of urban roads. AWSIM allows us to test HD map modifications in a controlled setting while maintaining real-world complexity.

HD Map and Editing Tools. For the HD map, we utilize an open-source vector map of Nishi-Shinjuku, which includes detailed representations of road geometry, lane structures, traffic signals, and pedestrian crossings. We modify this map using *Vector Map Builder* (VMB) [26], a tool that enables precise lanelet editing. Specifically, we manipulate lanelet widths within an intersection, forcing the AV's planning module to generate a new trajectory based on the altered lane structure.

4.3 Attack Implementation and Evaluation Methodology

Our tempering attack on lanes consists of three main steps:

Step 1: Lane Information Modification. The adversary manipulates the HD map by expanding the width of a selected lanelet at an intersection. The lanelet, a fundamental unit of Lanelet2, represents a segment of the road between two lane boundaries. Lanelet2 is a C++ library widely used in the *Robot Operating System* (ROS) ecosystem for handling high-definition maps in autonomous driving. It provides a flexible and extensible framework for representing and modifying road geometries, ensuring consistency across modules such as perception, planning, and control.

In this experiment, a lanelet measuring 57 meters in length is modified to induce planning deviations. The modification is performed using *Vector Map Builder* (VMB), a tool for editing Lanelet2-compatible maps. The tampered map is then loaded into Autoware, where the autonomous vehicle (AV) generates

Table 2: HD Map Tampering Attacks and Their Impact.

Tampered Feature	Affected Map Layer	Target Module	Impact
Lane (Addition, Deletion, Deformation)	Geometric Map, Road Connectivity	Mission Planner, Localization	Forcing AVs into incorrect lanes
Centerline (Modification)	Geometric Map, Road Connectivity	Mission Planner, Localization	Lane deviation, unsafe maneuvers
Road Attributes (Speed Limit, Curvature)	Semantic Map	Behavior Planner	Unsafe acceleration, braking
Road Attributes (Vector Direction, ID)	Semantic Map	Mission Planner	Incorrect route generation
Traffic Signs (Position Change)	Geometric Map	Localization	Misleading navigation
Traffic Signs (Addition, Deletion, Content Change)	Semantic Map	Perception	Incorrect traffic rule interpretation
Traffic Information (Fake Data Injection)	Real-Time Map	Behavior Planner, Motion Planner	Routing AVs into congestion or hazards
Point Cloud (Data Manipulation)	Base Map	Localization, Behavior Planner	Disrupting localization accuracy

a new trajectory based on the altered lane information. This manipulation demonstrates how an adversary can exploit Lanelet2-based map dependencies in ROS to influence AV behavior.

Step 2: Trajectory Generation in Autoware. The modified HD map is loaded into Autoware, and the AV is instructed to navigate through the intersection. The initial position is set just before the modified lanelet, and the destination is placed beyond the intersection. The planning module generates a trajectory based on the altered lane information.

Step 3: Evaluation of Planning and Driving Behavior. We measure deviations in the AV's planned path by comparing trajectories generated under normal and tampered HD maps. Additionally, we analyze the actual driving behavior in the simulation, assessing whether the AV crosses lane boundaries or enters oncoming traffic.

Evaluation of Planning Deviation. To evaluate the impact of lanelet modifications on planning, we first determine a threshold for unsafe trajectory deviations based on road width and vehicle dimensions. Inspired by adversarial backdoor attack evaluations [21], which introduce hidden triggers to manipulate model behavior, we define a deviation threshold th based on the lane width w_l and vehicle width w_v. In our study, we investigate how small lanelet modifications can act similarly, subtly altering the AV's planned trajectory.

The modified lanelet has a width of approximately 3.0 meters, consistent with standard urban road widths in Nishi-Shinjuku, the area covered by the HD map we use. The autonomous vehicle used in our simulation, a Lexus RX450h 2015, has a width of 1.895 meters ($w_v = 1.895$ m). To ensure that the vehicle remains within its designated lane and does not encroach on adjacent lanes, we set an allowable deviation threshold of $th = 0.5$ m.

Next, we systematically expand the lanelet width to $w_l = 3.5, 4.0, 4.5, 5.0$ meters (Figure 3) and evaluate its impact on the planned trajectory. Specifically, we measure the deviation δ at five key locations: starting point of the lanelet(0 m), its endpoint (57 m), two pedestrian crosswalks, and within the intersection. To quantify these deviations, we use the Fréchet distance d_F, a widely used metric for trajectory similarity evaluation [11, 14], which is defined as:

$$d_F(P,Q) = \inf_{\alpha,\beta} \max_{t\in[0,1]} \|P(\alpha(t)) - Q(\beta(t))\|, \qquad (1)$$

where $\|\cdot\|$ represents the Euclidean distance, and the trajectories are given as coordinate sequences $P = \{p_1, p_2, \ldots, p_n\}$ and $Q = \{q_1, q_2, \ldots, q_m\}$, with α and β being reparameterization functions [18].

We then analyze the maximum deviation for each lane width expansion and examine its correlation with the extent of lanelet modifications.

End-to-End Evaluation of Simulation-Based Driving Behavior. To analyze the full impact of lanelet modifications, we assess how the simulated autonomous vehicle behaves when following a trajectory planned using the tampered HD map. Similar to our planning evaluation, we compute the Fréchet distance between the vehicle's actual driving trajectory and the original trajectory under unmodified conditions. This allows us to assess how lanelet modifications influence real-time behavior of the vehicle.

Furthermore, we investigate the impact of adding a centerline to the modified lanelet. In this scenario, the lanelet is extended into the opposing lane, and a centerline is placed to guide the AV into the altered path. The centerline is designed to ensure a smooth transition between lanelets, preventing abrupt stops along the trajectory.

4.4 Results

Impact on Planning. First, we present the results of the planning attack. The lanelet width w_l was incrementally expanded, and the deviation δ from the planned trajectory was measured

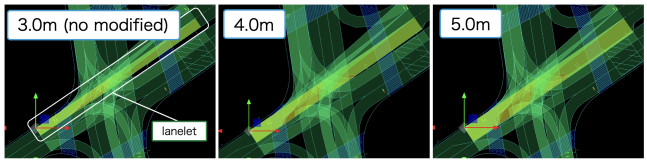

Figure 3: Visualization of lanelet expansion.

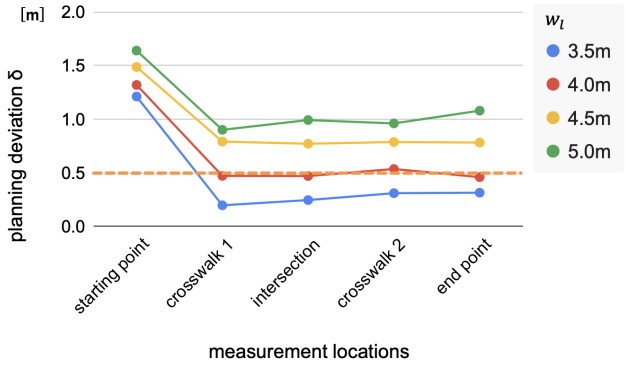

Figure 4: Relationship between measurement locations and planning deviation δ.

at five locations: the start of the lanelet (0 m), its endpoint (57 m), two crosswalks, and inside the intersection. Figure 4 shows the results. We also computed the Fréchet distance between the original planning trajectory and the one generated under each expanded lane width, summarized in Table 3.

The results indicate that when the lane width w_l exceeds 4.0 meters, the planning deviation δ surpasses the threshold $th = 0.5$ meters. This suggests that the planned trajectory shifts in proportion to the lanelet expansion. Consequently, attackers could induce lane deviation in an autonomous vehicle with minimal modifications to the lanelet width.

Impact on End-to-End Driving Behavior. Next, we analyze the effect of planning modifications on the AV's actual driving trajectory in simulation. We first consider the case without a centerline modification. For each lanelet width w_l, we compare the driven trajectory based on the modified planning path with the trajectory under the unaltered lanelet using Fréchet distance, as shown in Table 3. The reported values are the average Fréchet distance over three simulation runs. In the case of $w_l = 5.0$, the vehicle stopped prematurely due to an unrealistic planned trajectory, and thus this case was excluded from Fréchet distance calculations. The relationship between lanelet width expansion and both planning and driving trajectory deviations is summarized in Figure 5. Similar to the planning deviations, the vehicle trajectory deviations increased as w_l increased. In addition, when w_l reached 5.0 meters, the vehicle failed to reach its intended destination and stopped before completing the route. This suggests that abrupt changes in lane width can lead to infeasible planning solutions. If an attacker aims to manipulate an AV's movement to a target

Table 3: Lanelet width expansion and corresponding Fréchet distances d_Fp, d_Fe.

w_l [m]	d_Fp [m]	d_Fe [m]
3.5	0.5231	0.6049
4.0	0.7590	0.8419
4.5	1.0330	1.0965
5.0	1.2719	–

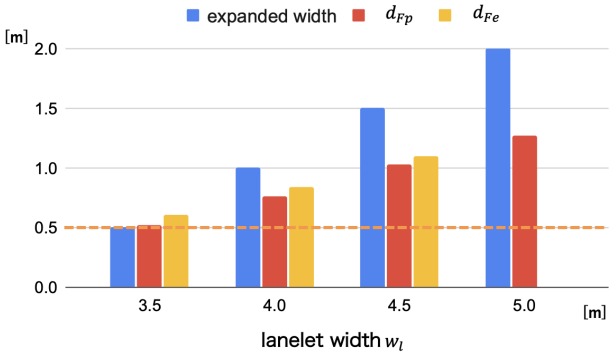

Figure 5: Relationship between expanded lane width w_l and Fréchet distance.

location, simply changing the lanelet width is not enough. Instead, a gradual change in lane width must be applied to maintain plausible planning results.

We then evaluate the case where a centerline is added to the modified lanelet. The resulting vehicle trajectories are shown in Figure 6. This result indicates that once a centerline is introduced, the AV strictly adheres to it, meaning that attackers can use centerline modifications to direct the AV along specific paths. This demonstrates the need for robust validation mechanisms to detect and prevent adversarial HD map modifications.

5 Discussion

Limitations. This study was conducted in a simulation environment using Autoware and AWSIM, which are widely used in both research and commercial autonomous vehicle development. While these platforms provide a strong foundation, real-world conditions introduce additional complexities that may not be fully captured. Our experiments relied on Autoware vector maps, but different AV systems may use different HD map formats, potentially affecting generalizability. Moreover, the evaluation focused on intersection scenarios, leaving the impact of HD map manipulation in other environments as an open question. Further studies should testing additional platforms and real-world deployments, as well as explore additional attack vectors such as traffic sign manipulation and the balance between reliance on HD maps and real-time lane detection as a defense strategy. A structured risk assessment

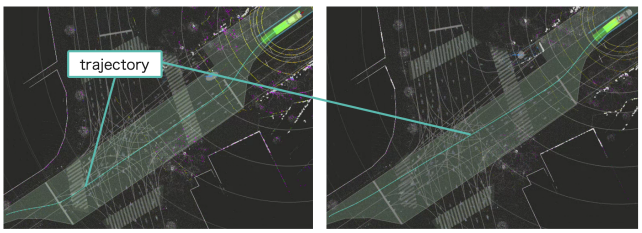

Figure 6: Comparison of driving trajectories with and without a centerline (left: no centerline, right: with centerline).

approach should also be considered to better understand the potential impact of HD map manipulation on autonomous driving systems.

Simulation versus Reality. While our experiments were conducted using Autoware and AWSIM in a controlled simulation environment, real-world deployments of autonomous vehicles face additional challenges. In practice, autonomous systems rely on sensor fusion from LiDAR, cameras, and radar to interpret road conditions. This multi-modal perception may compensate for some map inaccuracies or tampering attempts. However, under adverse conditions such as snowfall, dust accumulation, faded or obscured lane markings, or intersections with discontinuous paint, sensor visibility becomes limited, increasing reliance on HD map data. Consequently, tampered or outdated HD maps could have a more severe impact on vehicle behavior in these contexts than observed in simulation. This underscores the importance of evaluating map manipulation not only in idealized environments but also under realistic operating conditions, which we leave as future work.

Generalisability Across Stacks. Tampering vulnerabilities we identified are likely to affect a wide range of autonomous driving systems that depend on HD maps. Our findings were derived from experiments using Autoware vector maps, but many autonomous driving systems—such as Apollo [1]—implement similar core functionalities that rely on HD maps, particularly for tasks like path planning and behavior prediction. Regardless of the specific map format (e.g., OSM-based vector maps, OpenDRIVE [2]), lane-level information and routing structures are consistent components across platforms. As HD map standardization progresses, it is likely that map tampering vulnerabilities will manifest similarly in other systems that depend on such data. However, the generalizability of these findings should be approached with caution when considering systems that do not rely heavily on HD maps. For instance, Tesla vehicles primarily use real-time perception rather than detailed map priors, although they still incorporate basic map data for certain navigation functions. Therefore, our results are most applicable to map-dependent stacks and should be interpreted accordingly when extending to systems with different architectural assumptions.

Countermeasures. To ensure the integrity of HD maps, it is essential that developers carefully verify map data using available tools. For example, Vector Map Builder can be used to check road connectivity and identify inconsistencies prior to deployment. Strengthening HD map update mechanisms is also critical. Implementing cryptographic verification and digital signatures for map data updates can help authenticate sources and prevent unauthorized changes.

Ethical and Safety Considerations. The threats identified in this study are not specific to any particular product or system but stem from the structural characteristics of HD maps and the general framework of autonomous driving technology. Going forward, we plan to share our findings with map data providers and strengthen security measures.

6 Conclusion

This study investigates the security risks of HD map tampering attacks, where adversaries manipulate map data to mislead autonomous vehicles (AVs), creating systemic threats beyond sensor-based adversarial attacks. We categorized HD map tampering threats based on targeted map elements and their impact on AV operations, providing a structured understanding of potential vulnerabilities. Through simulation experiments with Autoware and AWSIM, we showed that even small modifications to lane widths can significantly alter AV trajectories, causing unsafe behavior. To mitigate these risks, we discussed the importance of rigorous HD map verification and enhanced update mechanisms. Future work should explore additional attack vectors, refine defensive strategies that balance HD map reliance with real-time perception, and establish a structured risk assessment framework. This study contributes to the AV security by both classifying HD map manipulation threats and demonstrating the need for robust defenses.

Acknowledgments

This work was partially supported by JSPS KAKENHI 22H00519 and JST CREST JPMJCR23M4.

References

[1] apollo. https://developer.apollo.auto/.

[2] opendrive. https://www.asam.net/standards/detail/opendrive/.

[3] OpenStreetMap. https://www.openstreetmap.org.

[4] WAYMO. https://waymo.com/.

[5] Ghadeer Abdelkader, Taghreed Alghamdi, Khalid Elgazzar, and Alaa Khamis. HD Maps for Connected and Automated Vehicles: Enabling Technologies and Future Directions. In *Proceedings of the 2023 IEEE International Conference on Smart Mobility (SM)*, 2023.

[6] Kaleab Taye Asrat and Hyung-J Cho. A Comprehensive Survey on High-Definition Map Generation and Maintenance. *ISPRS International Journal of Geo-Information*, 13(7), 2024.

[7] SHI Bo. OpenHDMap. https://github.com/bhsphd/OpenHDMap-1.

[8] Pino Bonetti. HERE introduces HD Live Map to show the path to highly automated driving. https://www.here.com/learn/blog/here-introduces-hd-live-map-to-show-the-path-to-highly-automated-driving.

[9] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Z. Morley Mao. Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 2267–2281, 2019.

[10] Baodong Chen, Wei Wang, Pascal Sikorski, and Ting Zhu. Adversary is on the Road: Attacks on Visual SLAM using Unnoticeable Adversarial Patch. In *33rd USENIX Security Symposium (USENIX Security)*, pages 6345–6362, 2024.

[11] Ming Cheng, Bowen Zhang, Ziyu Wang, Ziyi Zhou, Weiqi Feng, Yi Lyu, and Xingjian Diao. VeTraSS: Vehicle Trajectory Similarity Search Through Graph Modeling and Representation Learning, 2024.

[12] AUTOMOTIVE EDGE COMPUTING CONSORTIUM. Operational behavior of a high definition map application white paper, 2020.

[13] Gamal Elghazaly, Raphaël Frank, Scott Harvey, and Stefan Safko. High-definition maps: Comprehensive survey, challenges, and future perspectives. *IEEE Open Journal of Intelligent Transportation Systems*, 4:527–550, 2023.

[14] Maurice Frechet. Sur quelques points du calcul fonctionnel. 1906.

[15] FZI Research Center for Information Technology. Lanelet2. https://github.com/fzi-forschungszentrum-informatik/Lanelet2.

[16] CNN Jamiel Lynch. Family sues google alleging its maps app led father to drive off collapsed bridge to his death, attorneys say. https://edition.cnn.com/2023/09/21/us/father-death-google-gps-drive-off-bridge-lawsuit-north-carolina/index.html.

[17] Shinpei Kato, Shota Tokunaga, Yuya Maruyama, Seiya Maeda, Manato Hirabayashi, Yuki Kitsukawa, Abraham Monrroy, Tomohito Ando, Yusuke Fujii, and Takuya Azumi. Autoware on board: Enabling autonomous vehicles with embedded systems. In *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS)*, pages 287–296. IEEE, 2018.

[18] Johann Laconte, Abderrahim Kasmi, Romuald Aufrère, Maxime Vaidis, and Roland Chapuis. A survey of localization methods for autonomous vehicles in highway scenarios. *Sensors, 2021, 22 (1)*.

[19] Qian Luo, Yurui Cao, Jiajia Liu, and Abderrahim Benslimane. Localization and Navigation in Autonomous Driving: Threats and Countermeasures. *IEEE Wireless Communications*, 26(4), 2019.

[20] K. Massow, B. Kwella, N. Pfeifer, F. Häusler, J. Pontow, I. Radusch, J. Hipp, F. Dölitzscher, and M. Haueis. Deriving HD maps for highly automated driving from vehicular probe data. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, 2016.

[21] Mozhgan Pourkeshavarz, Mohammad Sabokrou, and Amir Rasouli. Adversarial backdoor attack by naturalistic data poisoning on trajectory prediction in autonomous driving. In *IEEE/CVF Conference on CVPR*, 2024.

[22] Takami Sato, Yuki Hayakawa, Ryo Suzuki, Yohsuke Shiiki, Kentaro Yoshioka, and Qi Alfred Chen. LiDAR Spoofing Meets the New-Gen: Capability Improvements, Broken Assumptions, and New Attack Strategies. In *Network and Distributed System Security Symposium (NDSS)*, 2024.

[23] Takami Sato, Junjie Shen, Ningfei Wang, Yunhan Jia, Xue Lin, and Qi Alfred Chen. Dirty Road Can Attack: Security of Deep Learning based Automated Lane Centering under Physical-World Attack. In *30th USENIX Security Symposium (USENIX Security)*, pages 3309–3326, 2021.

[24] Takami Sato, Ningfei Wang, Yueqiang Cheng, and Qi Alfred Chen. A Cross-Verification Approach with Publicly Available Map for Detecting Off-Road Attacks against Lane Detection Systems. In *Proceedings of the NDSS Workshop on Vehicle Security and Privacy (VehicleSec)*, 2024.

[25] TIER IV. AWSIM. https://github.com/tier4/AWSIM.

[26] TIER IV. Vector Map Builder. https://tools.tier4.jp/feature/vector_map_builder_ll2.