

Poster: Security Analysis of Multi-Sensor Fusion based Localization in Autonomous Vehicles

Junjie Shen*, Jun Yeon Won*, Shinan Liu[†], Qi Alfred Chen*, Alexander Veidenbaum*

*University of California, Irvine; [†]University of Electronic Science and Technology of China

{junjies1, junyeonw, alfchen, alex.veidenbaum}@uci.edu; liushinan63@gmail.com

Abstract—Precise and robust localization is highly critical for making correct and safe driving decisions in Autonomous Vehicles (AVs). To achieve this goal, they are usually equipped with redundant and complementary sensors, e.g., LiDAR, GPS, and IMU, and use Multi-Sensor Fusion (MSF) algorithms to combine the observations. However, it is largely unclear how robust these MSF algorithms are in presence of practical sensor attacks such as GPS spoofing. In this poster, we share our recent progress in performing the first security analysis of representative MSF designs in AV systems under GPS spoofing attack.

I. INTRODUCTION

In recent years, Autonomous Vehicles (AVs) have started to enter our daily life. Waymo launched a self-driving taxi service in Phoenix, Arizona; Baidu collaborated with a vehicle manufacture to deploy their open-source AV platform Apollo [1] on shuttle buses. Within the AV system, localization, which estimates the real-time location of the vehicle, is one of the most important modules that are critical for making correct driving decisions. To improve the robustness, AV systems generally adopt a Multi-Sensor Fusion (MSF) design in the localization module [2], [3]. In the MSF, multiple localization sensors such as Inertial Measurement Units (IMU), GPS, and LiDAR are fused together using algorithms such as Kalman Filter (KF) to provide a robust location estimation [2].

While MSF can effectively increase the robustness for sensor noises, it is still unclear how robust it is under deliberate sensor attacks. In the context of AV localization, a classic but still practical sensor attack vector is GPS spoofing, which has been concretely demonstrated on various real-world systems such as drones, yacht, and smartphones. Unfortunately, due to the lack of cryptographic protection for civilian GPS usage, such spoofing is fundamentally difficult to be fully prevented today [4]. Considering the vital importance of AV localization to correct AV driving decision making and thus road safety, it is highly necessary to understand its security property under GPS spoofing.

In this work, we perform the first systematic security analysis of MSF-based AV localization under GPS spoofing. To perform the analysis, we first design synthetic scenarios to overcome the challenges in the handling of sensor noises and the lack of ground truth. Next, we formulate the analysis task as an optimization problem to understand the upper bound of attack capability, and then leverage the analysis insights to identify effective spoofing strategies. Our preliminary results show that a well-designed spoofing strategy is able to deviate the localization estimation of a representative MSF implementation by 2 meters in as short as 10 seconds. Our future work includes evaluating the discovered attack strategies on real-world traces and performing case studies on end-to-end driving

decisions in open-source real-world AV systems to understand the security implications on road safety.

II. THREAT MODEL

We assume a car-following model where the attacker is driving a car and following the victim AV with the same speed while launching the GPS spoofing. The practicality of such attack process has been demonstrated on real roads by previous work [5]. As the first step towards understanding the security properties, in this work we assume that the attacker has the access to the implementation of the MSF-based localization algorithm in the victim AV. This is possible when 1) the victim adopts a representative algorithm implementation that is publicly available, or 2) the attacker owns an AV of the same model as the victim and can reverse engineer it to analyze the algorithm binary. In §V, we discuss our plan to explore the possibility of launching the attack without such access.

III. ANALYSIS METHODOLOGY

In this section, we detail the key steps in our security analysis methodology.

Synthetic scenario. AV systems make decisions based on real-world sensor data. However, performing security analysis directly on real-world data is less effective since 1) the presence of sensor noises and biases makes it difficult to pinpoint the true cause of erroneous behaviors, and 2) real-world data does not include the ground truth localization to understand the actual attack effectiveness. Thus, in our analysis we create synthetic driving scenarios in which sensor data is *directly* reflecting the synthetic trajectory without any noise and bias.

As the first step, we model the simplest driving scenario: driving with a constant speed on a straight-line road. In this scenario, we set all fields in IMU data to zero except the gravity field. Since our analysis focuses on the security property of sensor fusion algorithm instead of individual non-fused localization sources, we replace the LiDAR locator with an ideal localization source (i.e., ground truth location) to minimize the effect of imperfect LiDAR locator designs on the analysis, as well as to reduce the modeling overhead. For the GPS data, we set the non-spoofed GPS location to the ground truth and set the standard deviation to the median value in real-world traces. During spoofing attack period in the analysis, we also set the GPS data standard deviation to the same value as the non-spoofed case.

Algorithm security analysis. The attack goal can be formulated as an optimization problem: finding the best spoofing strategy which can achieve the maximum deviation in MSF localization estimation. To efficiently solve the optimization problem, we apply *gradient ascent* to find the best spoofing




Fig. 1. Loss surface of the first spoofing GPS point.

parameters iteratively. Since we do not have the analytic form of the underlying optimization problem (i.e., the state equation of MSF), a numerical approximation of the gradient is used.

By solving the optimization problem in the synthetic scenario, we can identify the upper bound of the attack effectiveness, i.e., the maximum deviation of the localization estimation output a GPS spoofing attack can achieve. This is an upper bound since the optimization process has the knowledge of the localization module output, which is not accessible by the attacker during the actual attack time. Drawing insights from the analysis results, we then identify a set of spoofing strategies without such knowledge and analyze their effectiveness compared to the upper bound.

Attack evaluation and case studies. To understand the performance of the discovered attack strategies, we perform attack evaluation on real-world sensor traces, e.g., those released by Baidu Apollo or open datasets such as KITTI. In addition, to understand the security implications of these discovered attacks in AV driving scenarios, we perform case studies by launching these attacks during the operation of open-source AV systems such as Baidu Apollo in common real-world driving scenarios using simulation. Through these case studies, we demonstrate the attack impact on the end-to-end AV driving decision process and the potential damage on road safety.

IV. PRELIMINARY RESULTS

We use Baidu Apollo’s MSF implementation [2] as an example to perform the security analysis. The GPS spoofing attack will introduce one incorrect location per second since the GPS receiver operates at 1 Hz in Baidu Apollo. We refer to the spoofed GPS locations as spoofing points in the analysis.

Fig. 1 shows the loss surface of the first spoofing point, plotted by sampling fine-grained *spoofing parameters* (i.e., distances and degrees of the spoofed locations away from the ground truth trajectory). As shown, naive spoofing choices can only reach sub-optimal deviations. Also, the outlier detection commonly used in MSF is taking effect when the spoofing distance exceeds some certain threshold. We also find that the optimal parameters vary across different spoofing points, which indicates that a carefully designed spoofing strategy (i.e., spoofing parameters for the spoofing points) is required to reach the maximum deviation over a series of spoofing points.

Analysis results. The analysis results for the upper bound of the attack effectiveness on the first 10 spoofing points is shown in Fig. 2. We performed this analysis using two approaches: sampling and gradient ascent. In the sampling based approach, we enumerate different spoofing parameters for *each* spoofing point to get the best parameters. In the gradient ascent, the numerical approximation based optimization is used to improve the efficiency of parameter searching. As shown,




Fig. 2. Deviation upper bounds from the sampling based and the gradient based approaches. Two *naive* spoofing strategies are also shown in the figure.

the deviation upper bound at the 10-th point is close to 2 meters, which is already large enough to cause the victim AV to locate itself on a wrong traffic lane.

Fig. 2 also shows the deviations of two *naive* spoofing strategies: 1) fixed distance, which sets all the spoofing points away from the ground truth with a constant distance (2 meters in the lateral distance), and 2) fixed step size, where we gradually increase the distances of the spoofing points to the ground truth with a fixed step size (0.2 meter). As expected, the two strategies can barely cause any deviation. The reason is twofold: 1) naive spoofing strategies do not consider the change of optimal parameters across different spoofing points, and 2) the parameters have to be carefully chosen to avoid being detected by the outlier detection in Apollo’s MSF implementation. To achieve effective spoofing, we plan to analyze the reasons behind the high effectiveness found in the upper bound analysis and design spoofing strategies based on the insights. For example, as indicated in Fig. 2 one strategy may be increasing the spoofing distances exponentially, which can help effectively avoid the outlier detection.

V. CONCLUDING REMARKS AND FUTURE PLANS

In this work, we perform the first systematic security analysis of GPS spoofing on a representative MSF design. Our initial results demonstrate that a well-design GPS spoofing strategy is able to achieve large deviation in the localization estimation. Building upon these results, we plan to 1) relax the attack requirement by introducing noises in the spoofing parameters and standard deviation, and 2) evaluate the attack success rate on real-world sensor traces and perform case studies on the end-to-end driving decisions made in Baidu Apollo. We also plan to extend the threat model by removing the assumption that the attacker has access to the victim’s MSF implementation. For this, we plan to implement an MSF algorithm ourselves following a common design and study the transferability of its spoofing strategies to the MSF implementation in Baidu Apollo.

REFERENCES

- [1] Baidu, “Baidu Apollo,” <https://github.com/ApolloAuto/apollo>.
- [2] G. Wan, X. Yang, R. Cai, H. Li, Y. Zhou, H. Wang, and S. Song, “Robust and precise vehicle localization based on multi-sensor fusion in diverse city scenes,” in *ICRA’18*.
- [3] R. Karlsson and F. Gustafsson, “The future of automotive localization algorithms: Available, reliable, and scalable localization: Anywhere and anytime,” *IEEE signal processing magazine*, 2017.
- [4] A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, “Unmanned aircraft capture and control via GPS spoofing,” *Journal of Field Robotics*, 2014.
- [5] K. C. Zeng, S. Liu, Y. Shu, D. Wang, H. Li, Y. Dou, G. Wang, and Y. Yang, “All your GPS are belong to us: Towards stealthy manipulation of road navigation systems,” in *USENIX Security’18*.

Security Analysis of Multi-Sensor Fusion based Localization in Autonomous Vehicles

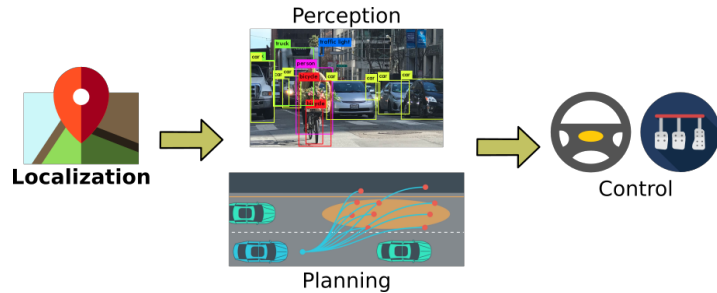
Junjie Shen, Jun Yeon Won, Shinan Liu¹, Qi Alfred Chen, Alexander Veidenbaum

University of California, Irvine, ¹University of Electronic Science and Technology of China

Autonomous Vehicle Localization


Autonomous Vehicle (AV) localization

- Real-time localization for making correct driving decisions
- One of the most **critical** modules: affects perception/planning/control



Multi-Sensor Fusion (MSF)

- Kalman Filter**: fuse location sensors, e.g. IMU, GPS, LiDAR, for robust localization
- Outlier detection**: reject anomalies in sensor data



GPS Spoofing and Threat Model

GPS spoofing attack

- Inject **falsified locations** in victim's GPS receiver
- A low-cost GPS spoofer can be **as cheap as \$225**
- Fundamentally hard to prevent** for civilian GPS receivers

Threat model

- Assume a **car-following model**: attacker drives at same speed as AV
- MSF implementation available**
 - AV adopts a representative MSF implementation that is publicly available (e.g., released by open-source AV platform like Baidu Apollo)
 - Or, attacker owns an AV of the same model with the victim and can reverse engineer it



Analysis Methodology

Synthetic scenario for effective analysis


- Security analysis on real-world sensor data is **less effective**
 - Sensor noises and biases** make it hard to pinpoint true causes for erroneous behaviors
 - No ground truth** for understanding attack effectiveness
- Create synthetic scenario: sensor data **directly reflects the trajectory**

Algorithm security analysis


- Attack goal**: find best spoofing strategy to achieve maximum deviation
- Identify the **upper bound of the attack effectiveness**
 - Numerical approximation based gradient ascent
- Draw insights of creating **effective spoofing strategies**

Attack evaluation and case studies

- Evaluate spoofing strategies using **real-world sensor traces**
- Case studies by launching these attacks in open-source AV systems



Preliminary Results




Experiment setup

- Target MSF binary in Baidu Apollo
- Synthetic scenario:
 - Constant speed on a straight line road

Loss surface of the 1st spoofing point

- Naive spoofing choice can only reach sub-optimal deviation
- Outlier detection: larger distance cannot reach larger deviation



Analysis results

- Well-designed spoofing strategies can cause **2-meter deviation in only 10s**
- Naive spoofing strategies can barely reach any deviation