A Multi Method Framework for GNSS Anomaly Detection in Vehicular Systems Using NMEA Data

Mathias Gerstner*, Tobias Reichel†, Sebastian Fischer*, Rudolf Hackenberg*
*Dept. Informatics and Mathematics, OTH Regensburg
Regensburg, Germany
e-mail: {mathias.gerstner, | sebastian.fischer, | rudolf.hackenberg}@oth-regensburg.de
†Central Office for Information Technology in the Security Sector
Munich, Germany
e-mail: tobias.reichel@zitis.bund.de

***Abstract*—Intended interference with satellite signals, including jamming and spoofing, has become increasingly common and poses a growing threat to safety-critical applications. In the automotive domain, this development creates a strong demand for reliable and lightweight detection mechanisms that can be deployed on standard vehicle platforms. This paper presents an Intrusion Detection System (IDS) designed to detect such attacks on common Global Navigation Satellite System (GNSS) receivers in road vehicle environments. The proposed system operates as a software solution and processes live, recorded, or simulated navigation data streams in real time. It integrates threshold checks, rule-based consistency analysis, and a Machine Learning (ML) algorithm together with a score fusion method to generate anomaly scores without requiring specialized hardware or labeled attack data. The system is adjusted and validated using real-world driving data combined with synthetic attack scenarios. Experimental results demonstrate reliable detection of diverse attack patterns while maintaining a low false positive rate under normal driving conditions, showing the suitability of the approach for automotive GNSS security monitoring.**

Keywords-gnss; nmea; autonomous driving; intrusion detection system; spoofing.

I. Introduction

Global Navigation Satellite System (GNSS), such as Global Positioning System (GPS) or Galileo, are fundamental components of modern navigation and localization applications. They are widely used in domains like aviation, maritime transportation, and automated driving to provide position, velocity, and precise timing information [1, pp. 10–14]. However, the open and largely unencrypted nature of civilian GNSS signals makes them vulnerable to cyberattacks such as jamming and spoofing, where counterfeit signals mislead receivers into reporting false navigation data [2]. As discussed in previous work, manipulated GNSS data can have significant safety, security, and forensic implications, particularly when navigation information is blindly trusted by systems [3].

To address these threats, GNSS-specific Intrusion Detection Systems (IDSs) have been proposed to monitor navigation data streams and identify anomalies indicative of attacks or system malfunctions [4]. Existing research has often focused on maritime environments or cooperative multi receiver setups exploiting spatial diversity. While these approaches can provide robust detection capabilities, they often require additional hardware or infrastructure support that limits their area of application [5]. In future autonomous driving scenarios, it will be important to have lightweight, configurable IDSs suitable for standard GNSS receivers commonly deployed in road vehicles.

This paper presents a framework designed to detect anomalies in National Marine Electronics Association (NMEA)-based navigation data streams for road vehicles. The system supports multiple input sources, including real-time data from physical GNSS receivers, replayed historical recordings, and simulated data generated by simulations, such as IPG CarMaker [6]. This design enables reproducible testing of attack scenarios and facilitates integration into real-time applications.

Anomaly detection is performed using parallel analysis methods. A threshold module evaluates navigation parameters against fixed bounds obtained from Exploratory Data Analysis (EDA) of real-world driving data. A rule-based component checks physical and logical consistency constraints, such as acceleration limits, velocity changes, or implausible short term deviations. In addition, a Machine Learning (ML)-based detector uses a Local Outlier Factor (LOF) algorithm to identify outliers in the navigation parameters without requiring labeled attack data.

The outputs of the individual detection modules are combined by a score fusion mechanism that computes a total anomaly score reflecting the presence and severity of detected irregularities. This modular architecture allows for a systematic evaluation of individual detection strategies while remaining suitable for deployment on automotive systems with limited computational resources.

Beyond standalone operation, the proposed IDS is intended to serve as a component of a forensic framework. Specifically, it is planned to be integrated into a Forensic Incident Recorder (FIR), currently under development, which extends the functionality of the Data Storage System for Automated Driving (DSSAD) [7]. Based on the severity level reported by the IDS, the FIR selectively triggers the recording of additional sensor data, such as speedometer or camera information, to balance storage constraints with forensic requirements in the event of suspected GNSS behavior [8].

The remainder of this paper is structured as follows. Section II reviews related work on GNSS intrusion detection. Section III describes the collection and exploratory analysis

of navigation data. Synthetic attack generation is presented in Section IV, while Section V details the system architecture and implementation. Section VI discusses the evaluation methodology and results. Section VII shows the system's capabilities and limitations, and Section VIII summarizes the findings and outlines directions for future work.

II. Related Work

Due to the increasing relevance of reliable satellite-based navigation, a growing amount of research addresses the detection of attacks on GNSS systems. Existing approaches include rule-based techniques as well as ML methods. However, a substantial part of the proposed solutions focuses on maritime applications or depends on specialized hardware. In contrast, lightweight and modular IDS architectures suitable for road vehicles remain underrepresented in the literature.

Amro et al. [9] propose an anomaly detection system for NMEA messages that relies on rules identifying violations of physical limits and timing constraints. ML methods were initially considered but discarded due to insufficient performance caused by limited training data. In contrast, the approach presented in this paper extends similar rule-based concepts by incorporating an unsupervised ML layer and fusing all detector outputs.

Boudehenn et al. [10] developed an IDS that operates on a Raspberry Pi and processes NMEA data. Their system employs a ML model to detect spoofing and jamming attacks. While computationally efficient for specific attack classes, it is tailored to maritime scenarios and relies on a single detection mechanism, limiting its adaptability to dynamic road vehicle environments.

Spravil et al. [11] introduced the MAritime Nmea based Anomaly detection (MANA), which detects GPS spoofing through the fusion of multiple consistency checks. However, MANA relies exclusively on predefined thresholds and rules, whereas the proposed architecture extends this concept by integrating anomaly detection through ML as a third detector.

Lemieszewski [12] proposed a physical consistency checking approach that compares the speed reported by a GNSS receiver with the vehicle's speedometer. This rule-based method demonstrated successful spoofing detection in real-world driving experiments but may fail for attacks that closely mimic realistic vehicle dynamics. In contrast, the IDS in this paper operates solely on GNSS data, enabling deployment without dependencies on additional vehicle sensors.

The developed system addresses these limitations by emphasizing the use of real-world driving data for analysis and detector calibration, complemented by controlled synthetic manipulations. By processing GNSS data in real time and integrating multiple detection strategies, the IDS provides a practical solution for automotive GNSS security monitoring without much overhead.

III. Exploratory Data Analysis

An EDA was conducted to find the statistical properties of the GNSS data and to identify feature correlations relevant for anomaly detection. The analysis was based on 50 historical driving recordings collected on a fixed suburban route with a length of 71 km. The measurement drives were conducted on the road B16 in Germany during a time window between March and December 2025. The receiver used is the u-blox ZED-F9R with the active antenna u-blox ANN-MB-00-00. All drives were performed under normal operating conditions without intentional interference, providing a baseline of expected GNSS behavior. The resulting insights were used to derive threshold values, define detection rules, and select suitable input features for the ML detection component.

Correlation analysis of all captured, unaltered data points revealed several strong linear dependencies among some GNSS features. In particular, the Dilution of Precision (DOP) metrics showed strong positive correlations. Similarly, position uncertainty parameters, including the major and minor axes of the error ellipse, as well as latitude- and longitude-related standard deviations, were highly correlated with each other and with the DOP metrics. These results were to be expected, but must be taken into account when selecting parameters [13].

In contrast, dynamic motion features such as vehicle speed and course exhibited only weak correlations with signal quality parameters. This indicates that motion behavior provides independent information that is not directly influenced by the satellite constellation quality. Consequently, these features are suitable for detecting anomalies characterized by implausible dynamics, such as unrealistic accelerations or abrupt direction changes, even when signal quality indicators appear nominal.

Feature correlations play a crucial role in the selection of inputs for the ML detector. Highly correlated features introduce redundancy into the feature space and can negatively affect the performance of ML algorithms [14]. To mitigate this, only weakly or moderately correlated features were selected for the LOF algorithm, ensuring a balanced representation of independent behavioral and signal related characteristics.

In addition to correlation analysis, descriptive statistics of all driving data were collected for all GNSS parameters. Mean values, standard deviations, and percentiles were calculated to characterize normal operating ranges. These statistics form the base for threshold detection and help to reduce false positives caused by normal environmental variability. The thresholds used in the IDS are based on the data provided by the mentioned specific hardware. Therefore, they have to be adjusted if another GNSS receiver and antenna are used.

IV. NMEA Data Manipulation

To enable systematic evaluation of the IDS, an offline GNSS anomaly generation framework was developed that operates directly on recorded NMEA data stored in Comma Separated Values (CSV) format. The internal processing workflow, including EDA and model training, is illustrated in Figure 1. By manipulating existing recordings instead of generating fully synthetic navigation data, realistic temporal behavior and noise characteristics of real GNSS measurements are preserved.

In addition to offline manipulation, controlled spoofing experiments were conducted in a shielding tent to study NMEA

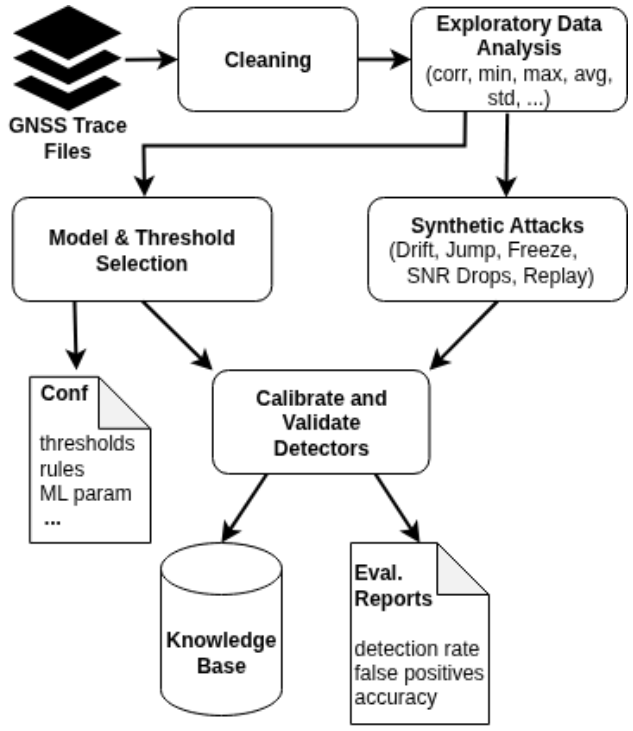

Figure 1. Overview of the EDA and model training process.

behavior under real spoofing conditions. Using a Software Defined Radio (SDR), counterfeit GNSS signals with incorrect satellite constellations were transmitted to a receiver. These experiments revealed effects such as temporary loss of valid navigation data, position freezes, and slow positional drifts. Shortly after the start of spoofing, NMEA sentences often contained empty or invalid fields before stable spoofed position and time information was created by the receiver. Since this behavior is comparatively easy to detect, offline manipulation assumes a near-ideal spoofing scenario with smooth transitions and no data gaps.

In the following, different effects of spoofing were applied to the NMEA data to emulate attacks and fault conditions. A position jump attack introduces a sudden spatial displacement by shifting all NMEA sentences containing geographic location by a constant, configurable offset starting at a defined time. The modified coordinates are converted back into NMEA sentences and the checksums are fitted to the new payload.

A drift attack models a drag-off spoofing scenario by gradually increasing the positional offset over a specified time interval. Each affected NMEA sentence is shifted by a time-dependent fraction of the maximum configured offset of 100 m, resulting in a smooth and continuous deviation from the true trajectory.

For a position freeze attack, all NMEA sentences containing position information are overwritten with a fixed location equal to the last valid position before the attack onset, while other fields, such as time, continue to evolve normally. This manipulation emulates a receiver that appears stationary despite ongoing movement.

Signal quality manipulation, as it may occur during jamming, is realized through Signal to Noise Ratio (SNR) degradation. Within a defined time window, NMEA GSV sentences are modified by identifying satellite SNR fields and scaling them by a configurable factor.

Finally, replay attacks are considered, representing a practical threat requiring minimal knowledge of receiver internals. In this scenario, authentic GNSS signals are captured using an SDR and later retransmitted. To increase realism and detection difficulty, an advanced replay strategy is assumed in which the captured signals are not simply replayed as captured. If the timestamps show an earlier time than the real time, such attacks are easy to detect because time cannot run backwards. Therefore, the timestamps in the replayed messages are set to the actual time, avoiding obvious temporal inconsistencies.

Together, these manipulation techniques provide a diverse set of attack scenarios that support evaluation of the detection capabilities of the GNSS IDS [15].

V. System Architecture

The developed IDS combines multiple complementary detection strategies to identify spoofing and jamming anomalies. Instead of relying on a single detection mechanism, the system integrates thresholds, rules, and ML approaches and aggregates their outputs in a score fusion module. An overview of the architecture is shown in Figure 2.

If NMEA data is produced by a simulation or a real GNSS receiver, it is usually embedded in other software. Therefore, a Message Queuing Telemetry Transport (MQTT) client is implemented to fetch this data from a broker, which acts as the interface between data producing components and the IDS itself. Reports or alerts generated by the IDS can also be sent back to other software via the MQTT interface.

In this study, we used GNSS trace files to train and validate the system's capabilities. Regardless of the NMEA data source, it is first stored in a queue for raw data. A processing module then extracts the relevant information from the NMEA records and places it into a separate queue for each detector. The detectors can then retrieve the current GNSS data from these queues and evaluate it using various methods. If a detector identifies an anomaly, it calculates a score depending on the severity of the anomaly. These scores are combined into a total anomaly score, which is mapped to a continuous severity scale ranging from 0 to 1, enabling graded alerting instead of binary decisions.

A. Threshold Detection

Threshold violation detection is the simplest and most computationally efficient component of the IDS. It relies on static upper or lower bounds derived from physical constraints and statistical properties observed during the EDA. Monitored parameters include geometry related quality indicators, position uncertainty estimates, satellite availability, SNR, and vehicle speed. The thresholds used in this work are summarized in

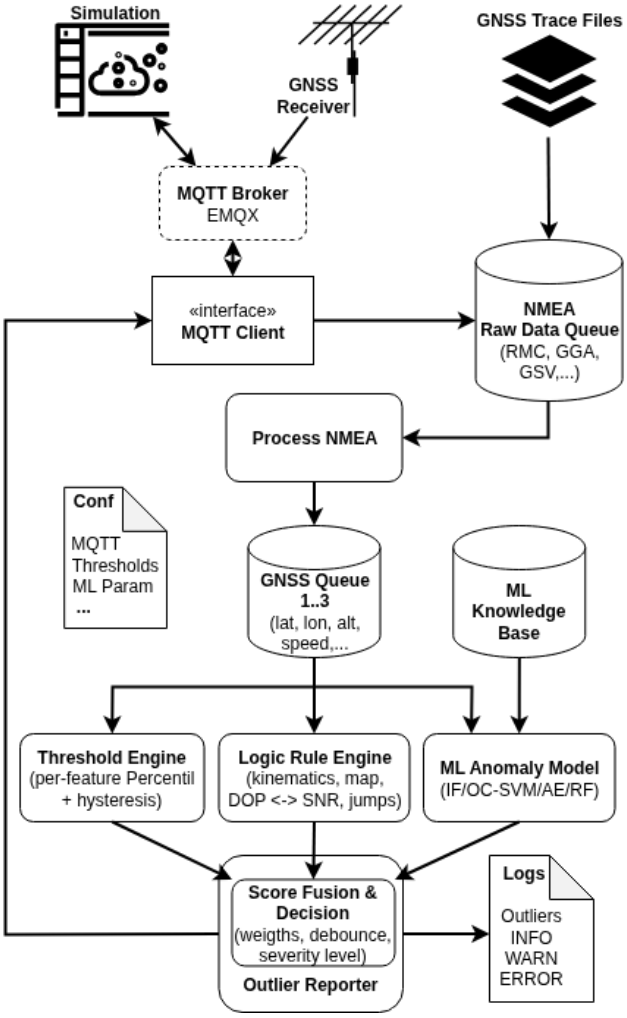

Figure 2. Internal System Architecture of the GNSS IDS.

Table I. The speed on the testing route is 100 km/h, so an alert limit of 120 km/h was chosen for an anomaly report.

For all metrics, additional overshoot thresholds are defined to capture extremely unrealistic values. The thresholds are defined by the extrema observed across all 50 drives, i.e., the maximum for some parameters and the minimum for others. Violations of these overshoot thresholds result in increased severity contributions, indicating a higher probability of malicious manipulation. Threshold checks are evaluated on a per epoch basis at 1 Hz, enabling immediate detection of very noticeable anomalies. While this approach lacks contextual awareness and may produce false positives in challenging environments such as urban canyons or tunnels, it provides an effective first line of defense.

B. Rule Detection

The rule-based detection module evaluates higher-level physical and logical consistency constraints within the GNSS data stream. Instead of analyzing individual parameters, it considers temporal and relational properties. The implemented rules are:

TABLE I. THRESHOLDS USED FOR ANOMALY DETECTION

Metric	Threshold	Overshoot Threshold
PDOP	> 2.03	> 22.54
HDOP	> 1.13	> 19.94
VDOP	> 1.71	> 22.01
RMS range	$> 45.0\,\mathrm{m}$	$> 135.0\,\mathrm{m}$
Std. major axis	$> 10.0\,\mathrm{m}$	$> 199.0\,\mathrm{m}$
Std. minor axis	$> 5.7\,\mathrm{m}$	$> 66.0\,\mathrm{m}$
Std. latitude	$> 4.0\,\mathrm{m}$	$> 109.0\,\mathrm{m}$
Std. longitude	$> 2.5\,\mathrm{m}$	$> 147.0\,\mathrm{m}$
Std. altitude	$> 6.7\,\mathrm{m}$	$> 500.0\,\mathrm{m}$
GPS satellites	< 9	$< 7 \vee > 17$
Mean SNR	$< 22.7\,\mathrm{dB}$	$< 5.0\,\mathrm{dB}$
SNR std. dev.	$< 6.2\,\mathrm{dB}$	$< 1.5\,\mathrm{dB}$
Max. speed	$> 120\,\mathrm{km/h}$	$> 160\,\mathrm{km/h}$

- Lack of incoming data for more than 5 seconds
- Acceleration (positive or negative) exceeding 14 m/s^2
- Position jump exceeding 80 m/s
- Yaw rate exceeding 60°/s (unrealistic steering maneuver)
- Jump between two epochs in the mean SNR over all satellites exceeding 12 dB

The threshold values were determined based on an EDA conducted over all 50 drives, which correspond to a total driving distance of 3,350 km.

This detector is effective at identifying structured inconsistencies that may not exceed individual thresholds but violate expected motion or signal behavior. However, like most deterministic approaches, it can be evaded by well-executed gradual attacks, such as slow drag-off spoofing, that remain within predefined limits over extended periods.

C. Machine Learning Detection

The third detector is based on ML. It uses a LOF algorithm that learns normal system behavior by analyzing density patterns in the parameters. The approach compares the current 20 seconds with the last 200 seconds and therefore does not need labeled attack data.

The detector operates on a rolling window of recent GNSS observations. Incoming epochs are evaluated at a frequency of 5 seconds. For each evaluation cycle, the most recent data is split into a training subset representing nominal behavior and a smaller test subset that is assessed for anomalies. All features are standardized based on the training data. This sliding window enables the detector to adapt to slowly varying environmental and operational conditions while remaining sensitive to short term deviations.

The selected feature set contains the delta of latitude and longitude (previous position to current location), speed, position DOP, major axis standard deviation, number of visible satellites, standard deviation of SNR, and the Root Mean Square (RMS) satellite range. Strongly correlated parameters are always represented by just one of them. By analyzing these parameters, the LOF detector captures dependencies that cannot be reliably addressed by threshold checks or fixed rules.

Anomalies are identified when test samples exhibit significantly lower local density compared to their neighborhood in the learned feature space [16]. To reduce sensitivity to minor fluctuations, only anomalies exceeding a predefined severity threshold of -5.0 are reported. Detected events are then sent to the fusion and reporting modules of the system.

The main limitation of this module is a moderate detection latency introduced by the window-based processing, which represents an acceptable trade-off for increased detection reliability.

D. Score Fusion

The final stage of the IDS architecture is a parallel operating score fusion module. Instead of binary decisions, the system aggregates, weights and normalizes the outputs of all detection layers. The fusion process produces a continuous total anomaly score, which is mapped to a severity value between 0 and 1. In the current implementation, an anomaly is reported when the total score exceeds 0.5, with all parameters configurable via a central configuration file.

This fused representation improves robustness against false positives from individual detectors and enables a balanced reaction strategy. Low severity anomalies trigger early warnings, while high confidence events supported by multiple detection layers escalate to critical alerts. For integration into the planned FIR, the continuous severity score will be mapped to discrete levels (LOW, MEDIUM, HIGH, and CRITICAL) to control severity dependent response and data recording behavior.

VI. Validation with Real World Data

The performance of the IDS is validated using real-world driving data. Two evaluation scenarios are considered: (i) detection performance under synthetic attack conditions and (ii) false positive behavior under nominal, attack-free operation.

A. Detection Performance

As described in Section IV, synthetic spoofing and jamming attacks are generated and injected into recorded NMEA files. These manipulated datasets are processed by the IDS to evaluate its ability to detect anomalous behavior.

Figure 3 shows the resulting anomaly scores for a representative drive containing different attacks. The IDS raises elevated anomaly scores during attack phases, while maintaining low scores during nominal operation outside the manipulated intervals.

A short detection latency is observed for slow drifts and location freeze attacks, as it is inherently difficult to distinguish between a legitimate vehicle stop and malicious manipulation. One false positive alert can be seen at the right border of the graphic, caused by the fast decrease of DOP values.

B. False Positive Analysis

In addition to detection capability, false positive behavior is a key performance metric for an IDS. To evaluate this aspect, the system was applied to a dataset containing 50 real-world driving runs on a rural road.

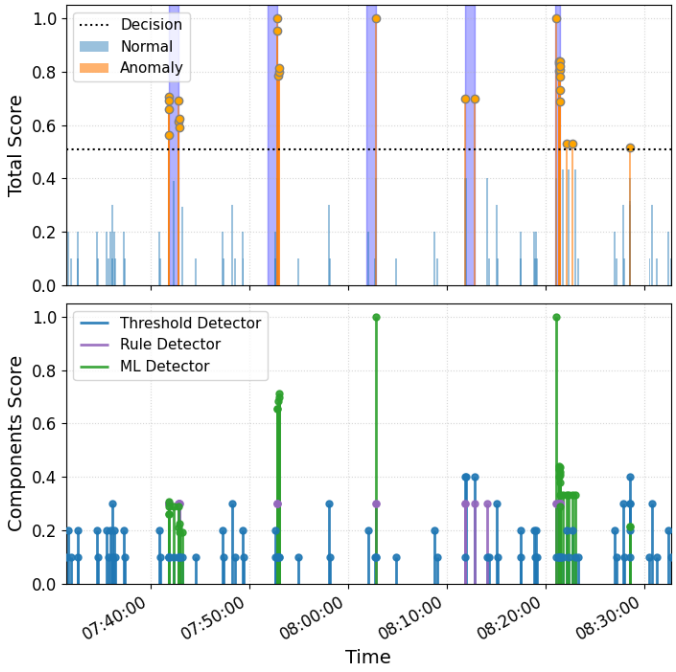

Figure 3. IDS anomaly total and per component scores during different spoofing and jamming attacks. Blue marks show the phases of attack. Orange marks show detected anomalies. From left to right the attacks are: position jump, slow drift attack, location freeze attack, jamming (SNR drop), replay attack.

Figure 4 presents the anomaly scores for the same route shown in Figure 3, but without injected attacks. The IDS maintains low anomaly scores throughout the drive, with only one single false alert observed, showing robustness during normal operation while remaining sensitive to abrupt inconsistencies.

For a systematic analysis, all driving runs were evaluated. Since the dataset contains no synthetic attacks, all detected anomaly events are treated as false positives. Consecutive anomaly alerts are counted as a single event until the system reports a normal state at least once. Each driving run comprises approximately 4000 GNSS epochs with 1 Hz sampling and a distance of 71 km.

TABLE II. Performance of the GNSS IDS evaluated on 3550 km of drives.

Metric	False positives per hour
Mean false positive rate	2.27
Median false positive rate	1.80
Maximum false positive rate	10.80

As shown in Table II, the system exhibits low false positive activity under normal conditions, with less than three events per hour on average. Analysis of runs with higher rates indicates that false positives are primarily caused by degradations in GNSS signal quality. The Table III contains the individual features and detectors that were most frequently involved in false-positive anomaly alerts. In total, 126 alerts were raised over all test drives. Within these, the threshold validation of the mean SNR was involved in 54 of them.

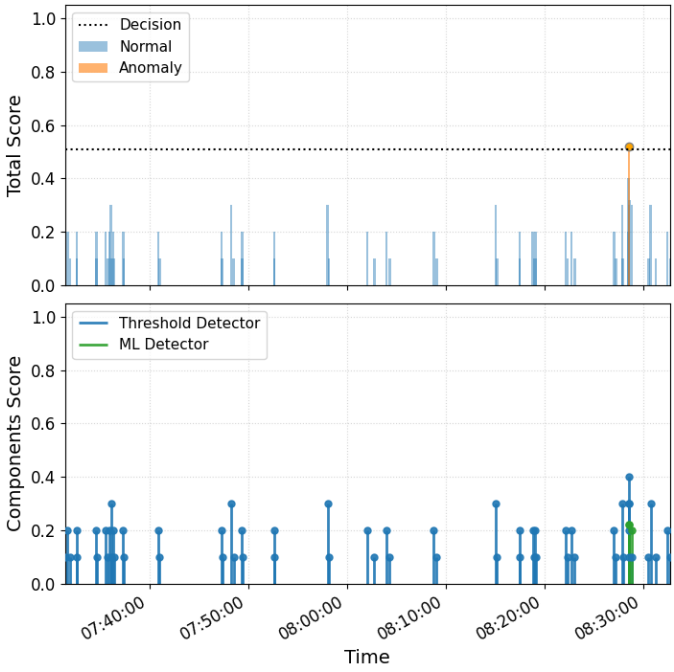

Figure 4. IDS anomaly scores under normal driving conditions without attacks.

TABLE III. MOST FREQUENT CAUSES OF FALSE POSITIVE EVENTS ACROSS ALL ATTACK FREE RUNS.

Detector and Feature	Occurrences
Threshold: mean SNR	54
LOF: RMS range	52
LOF: major axis position error	51
LOF: Position DOP	47
Threshold: RMS range	41

These results indicate that the system achieves a good balance between detection sensitivity and robustness against false alerts. While near perfect synthetic attacks are reliably detected, the false positive rate remains low under real-world driving conditions.

VII. DISCUSSION

Overall, the results show several practical challenges in the implementation of an IDS for navigation data. For example, rapid degradations of SNR or DOP values are not always associated with malicious interference. These effects frequently occur with specific road segments, such as loops, bridges, or forests, where signals are shadowed by obstacles. This behavior shows the importance of considering environmental influences when interpreting anomaly detections.

During the adjustment of the system's parameters, one task was the removal of the satellite SNR standard deviation from the rule-based detector. This feature is naturally correlated with the mean SNR which leads to simultaneous violations of both rules during signal degradations. As a result, these correlated reactions significantly increased the false positive rate.

By removing the standard deviation rule, the overall median false positive rate was reduced from 2.52 to 1.80 events per hour. However, this improvement came at the cost of reduced sensitivity in the jamming attack scenario. The overall severity score during this attack decreased from 1.0 to 0.7. This is still enough to raise an alert, but with a less clear statement. This illustrates the problematic trade-off between maximizing attack detection and maintaining an acceptable false positive rate. An alternative approach to handling SNR false positives could be the use of temporal gradient features. In particular, the rate of change of the SNR may provide a more discriminative indicator for abrupt interference events such as jamming while being less sensitive to naturally low but stable signal conditions.

The results further suggest that adding contextual information could help reduce false positives. For example, the use of digital maps with known influencing objects such as bridges, dense vegetation, or urban structures could allow the IDS to adapt its sensitivity locally. Such context-aware mechanisms could reduce alerts caused by predictable, non-malicious signal degradations without globally lowering detection rates.

It is also noteworthy that conventional performance metrics, such as the F1 score, were not applied in this study. This decision was made because it typically requires reliable ground truth for attacks, which is not available for real-world driving data. Instead, the false positive rate per fixed driving distance under normal conditions was chosen as the primary robustness metric. This approach provides more meaningful insight into system behavior in a real-world environment where attacks are rare and unknown, but too many false alerts can significantly limit practical usability.

VIII. CONCLUSION AND FUTURE WORK

The presented anomaly detection system for road vehicle environments operates on standard NMEA data. While the detection of anomalies works well, there are still a small number of false positives in data without attacks. Therefore, future work will focus on incorporating contextual information, such as additional vehicle dynamics, to verify plausible motion. Another promising concept is the integration of map and environmental data to further reduce these false positives and improve robustness against highly sophisticated spoofing attacks.

ACKNOWLEDGMENT

This work was supported by the project 'Digital Forensics in IT Systems (DiForIT)', funded by the German Federal Ministry for Economic Affairs and Climate Action (BMWK).

REFERENCES

[1] E. D. Kaplan and C. J. Hegarty, Eds., *Understanding GPS/GNSS: Principles and Applications*, 3rd ed. Boston, MA, USA: Artech House, 2017.

[2] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner, "Assessing the spoofing threat: Development of a portable GPS civilian spoofer", in *Proc. ION GNSS Conf.*, 2008.

[3] T. Reichel et al., "A forensic analysis of GNSS spoofing attacks on autonomous vehicles", in *Proc. 16th Int. Conf. Cloud Computing, GRIDs, and Virtualization, Valencia, Spain*, Apr. 2025, pp. 32–39.

[4] A. Mohanty and G. Gao, "A survey of machine learning techniques for improving global navigation satellite systems", *EURASIP J. Adv. Signal Process.*, vol. 2024, no. 1, p. 73, 2024. DOI: 10.1186/s13634-024-01167-7.

[5] M. L. Psiaki and T. E. Humphreys, "GNSS spoofing and detection", *Proc. IEEE*, vol. 104, no. 6, pp. 1258–1270, Jun. 2016. DOI: 10.1109/JPROC.2016.2526658.

[6] IPG Automotive GmbH, *CarMaker – Vehicle dynamics simulation*, [Online]. Available: https://www.ipg-automotive.com/solutions/product-portfolio/carmaker. Accessed: Feb. 17, 2026.

[7] United Nations Economic Commission for Europe (UNECE), "Guidance on data storage system for automated driving (DSSAD) and EDR", UNECE WP.29 GRVA, Tech. Rep., 2025, [Online]. Available: https://unece.org/sites/default/files/2025-06/GRVA-22-23e.pdf. Accessed: Dec. 18, 2025.

[8] K. Dolos et al., "Forensic readiness for autonomous mobility: The forensic incident recorder and information system concept", *Forensic Sci. Int.: Digit. Investig.*, vol. 56, Art. no. 302044, Mar. 2026. DOI: 10.1016/j.fsidi.2026.302044.

[9] A. Amro, A. Oruc, V. Gkioulos, and S. Katsikas, "Navigation data anomaly analysis and detection", *Information*, vol. 13, no. 3, p. 104, 2022. DOI: 10.3390/info13030104.

[10] C. Boudehenn, O. Jacq, M. Lannuzel, J.-C. Cexus, and A. Boudraa, "Navigation anomaly detection: An added value for maritime cyber situational awareness", in *Proc. Int. Conf. Cyber Situational Awareness, Data Analytics and Assessment (CyberSA)*, Jun. 2021, pp. 1–4. DOI: 10.1109/CyberSA52016.2021.9478189.

[11] J. Spravil, C. Hemminghaus, M. von Rechenberg, E. Padilla, and J. Bauer, "Detecting maritime GPS spoofing attacks based on NMEA sentence integrity monitoring", *J. Mar. Sci. Eng.*, vol. 11, no. 5, p. 928, 2023. DOI: 10.3390/jmse11050928.

[12] L. Lemieszewski, "Transport safety: GNSS spoofing detection using the single-antenna receiver and the speedometer of a vehicle", *Procedia Comput. Sci.*, vol. 207, pp. 3181–3188, 2022. DOI: 10.1016/j.procs.2022.09.375.

[13] W. Li et al., "GDOP and the CRB for positioning systems", *IEICE Trans. Fundam. Electron., Commun. Comput. Sci.*, vol. E100-A, no. 2, pp. 733–737, 2017. DOI: 10.1587/transfun.E100.A.733.

[14] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection", *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, 2003.

[15] Spirent Federal, "GNSS signal spoofing: How to evaluate the risks to safety-critical and liability-critical systems", Spirent Federal, Tech. Rep. DWP0014, Issue 1-02, 2020, [Online]. Available: https://spirentfederal.com/wp-content/uploads/Federal-DWP0014-Issue-1-02-GNSS-Signal-Spoofing-min.pdf, Accessed: Feb. 17, 2026.

[16] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: Identifying density-based local outliers", *SIGMOD Rec.*, vol. 29, no. 2, pp. 93–104, May 2000. DOI: 10.1145/335191.335388.