Anomaly detection of attacks on LIDAR based automotive perception systems

Jonas Schäufler
Department of Computer Science
HAW Hamburg
Hamburg, Germany
jonas.schaeufler@haw-hamburg.de

***Abstract*—As a result of advancements in bus technology and the development of functionality related to autonomous driving the connectivity of automotive vehicles and their components inside the vehicular network is increasing. This naturally leads to an expanding attack surface and creates the demand for defense mechanisms like intrusion detection systems. An integral task of an autonomous driving vehicle is the perception of the environment, based on which the vehicle controlling functions act according to a given situation. If an attacker is able to manipulate the percieved environment of the vehicle in a controlled way he would be able to prompt a calculated reaction of the vehicles controlling functions. As machine perception systems are designed to gather specific domain-dependant information, the task of detecting anomalous behaviour of these systems is also specific to their domain. Therefore it is necessary to create a domain specific solution to detect attacks on these systems on the application level. In this paper we present a concept for detecting the manipulation of application data of LIDAR based perception systems in an automotive context, and present evaluation results of a prototype implementaion.**

***Index Terms*—perception, automotive, LIDAR, anomaly detection, intrusion detection, binary bayes filter**

I. Introduction

The continouos development of driver assistance functions and an increasing demand of autonomously operating vehicles, as well as the required connectivity between vehicles and their components, lead to new challenges in terms of keeping the vehicles systems secure. The impact of attacks on cyber-physical systems like passenger vehicles is substantial as they may cause immediate physical harm to the people in the systems vicinity. The threat of an attacker directly accessing the vehicles controlling functions aiming to take control of the vehicle has been proven to be real and has been addressed in literature in various forms. Another way for an attacker to influence the vehicles behaviour is the manipulation of its percieved environment. If the attacker could manipulate environmental information in a controlled manner he would be able to prompt a calculated reaction of the components which are driving the vehicles actors. A common use case of machine perception in the automotive domain is to detect other traffic participants, required for autonomous driving and driver assistant functions like adaptive cruise control, lane change prediction and emergency braking. Assuming a Man-in-the-Middle scenario and the ability of an attacker to intercept and manipulate network traffic, an attacker could modify the location or speed of a detected object to trigger evasive maneuvers or the emergency brake. In this paper we address the problem of detecting these anomalies resulting from malicious manipulation of the machine perception interface. Based on a proposed system architecture for domain anomaly detection in machine perception [1] and state of the art anomaly detection mechanisms [2] the main goal of this paper is to develop the concepts needed to adapt this methods and the proposed architecture for the automotive domain focusing on the use case of object detection. In Section II we present related work and state of the art anomaly detection techniques. Recently LIDAR *(Light Detection and Ranging)* sensors are becoming the method of choice in the automotive industry to serve as input for machine perception algorithms, as they have higher resolution than RADAR sensors and are therefore able to detect much smaller objects. In Section III-B we examine the interface of such machine perception algorithms used for object detection in the automotive domain, present the possibilities of an attack, and formulate the problem statement. In Section IV-A we present our concept of calculating the probability of an anomalous prediction using a binary Bayes filter also used in occupancy grid mapping [3] and discuss the measurement model used for this approach. In Section V we present results of the evaluation of a prototype implementation and finally Section VI concludes the paper, presenting possible ways to improve upon our solution and gives an overview on future research.

II. Related Work

A. Anomaly detection techniques

Anomaly detection techniques have been studied extensively in literature and has been the topic in multiple surveys, articles and books. While most existing surveys heavily focus on a specific application domain or research area Chandola et al published a report [4] providing a structured and broad overview of the research on anomaly detection techniques spanning multiple research areas and application domains, including: *Cyber-Intrusion Detection*, *Fraud Detection*, *Medical Anomaly Detection*, *Industrial Damage Detection*, *Image Processing*, *Textual Anomaly Detection* and *Sensor Networks*. Generally techniques can be categorized into six different categories as shown in table I.

TABLE I
CATEGORIZATION OF ANOMALY DETECTION TECHNIQUES

Category	Technique
Classification Based	Neural Networks Based
	Bayesian Networks Based
	Support Vector Machines Based
	Rule Based
Clustering Based	K-means
	DBSCAN
	ROCK
	SNN clustering
	FindOut
Nearest Neighbor Based	distance to kTH Nearest Neighbor
	Relative Density Nearest Neighbor
Statistical	Gaussian Model Based
	Regression Model Based
	Mixture of Parametric Distributions Based
	Histogram Based
	Kernel Function Based
Information Theroetic	Kolomogorov Complexity
	relative entropy
	entropy
Spectral	Principal Component Analysis
	Compact Matrix Decomposition

Anomaly detection problems have multiple characteristics that vary from instance to instance and are not confined to any application domain, ultimately justifying the need for multiple different techniques. These characteristics include:

a) Input Data: The nature of input data is a key aspect of any anomaly detection technique. Generally the input is a set of data instances, consisting of one (*univariate*) or more (*multivariate*) attributes. Depending on the type and nature of attributes different techniques may be applicable to the specific anomaly detection problem. While the amount and types (*binary*, *categorical* or *continuous*) of attributes is one aspect of the inputs nature, another aspect is the relation between data instances. Relationships can be of sequential, spatial or graph based nature, or data instances may have no relationship amongst themselves, in which case they are referred to as record or point data. It is also possible for data instances to have mixed relationships, for example climate data, which can be considered *spatio-temporal* as instances are related to its neighbors and include the sequential component of time.

b) Type of Anomaly: Another important aspect is the nature of the anomaly to be detected. If an anomaly is deemed a data instance beeing an outlier regarding the set of other data instances, it is considered a *Point Anomaly*. If the data instance is considered anomalous only under a specific condition, but not otherwise, it is referred to as *Contextual Anomaly*. The context needs to be defined for a specific problem instance, and is generally given by or derived from a subset of attributes, therefore for *Contextual Anomalies* the data instance attributes are divided into *contextual* and *behavioral* attributes. If the data instances are related, an anomaly can also consist of a collection of data instances, of which a single instance must not be anomalous on its own. These anomalies are called *Collective Anomalies*. Depending on the data available a point or collective anomaly problem can also be transformed to a contextual anomaly problem.

c) Data Labels: The quality and availability of labeled data is another characteristic of an anomaly detection problem. Some techniques require labeled training data sets denoting the anomalous and normal data instances. This mode of operation is called *Supervised anomaly detection*. If a technique operates on data sets containing only labels for the non-anomalous instances it is considered *Semi-Supervised*. *Unsupervised* techniques do not require training data and rely on the assumption that anomalies are far more rare than normal instances.

d) Output of Anomaly Detection: Another aspect of an anomaly detection technique is the type of output it will produce. The output can either be a *label* or a *score* which are assigned to a data instance. While the *score* denotes the degree to which that instances is considered an anomaly, *labels* only define if an instance is anormalous or not.

B. Machine perception

In comparison to the topic of anomaly detection in general the application of techniques in the domain of machine perception is a relatively new application domain. As machine perception systems are designed to deliver a very specific functionality (like object detection) and the wide variety of different sensors that may be used in such a system require each technique to be adapted to the specific application and tailored to its use case. *Machine Perception* could be considered a subdomain of the *Sensor Network* domain and shares the same challenges that need to be overcome: *resource constraints*, *online processing*, *noise*, *missing data*, *distributed data sources* [4]. In addition *machine perception* algorithms operate on a higher level of abstraction and data of higher dimensions than raw sensor data processing algorithms. Having multiple levels of data representation (i.e. raw sensor data, clusters, features, objects) in a processing pipeline of machine perception algorithms leads to a new type of anomaly that can occur in such systems, where the interpretation of data on the different levels are incongruent. This type of anomaly is called a *compound anomaly*. Detecting incongruent events in multiple levels of data representation has been applied to novelty and rare behaviour detection in video analysis [5] [6] [7] and based on this pioneering work in detecting anomalous observations in perception systems Kittler et al aimed to develop a comprehensive framework architecture identifiying the mechanisms needed for determining the true nature of an anomaly in such systems. [1] Their system architecture incoorporates components for *observation anomaly detection*, *reject option detection*, *incongruence detection* and *sensor data quality gauging*. They emphasise that the detection of an anomaly and its qualification cannot be accomplished without these sources of information, but do not detail the actual anomaly analysis processes as they are problem specific.

III. Problem statement

To formulate the anomaly detection problem for this work, a basic architecture is shown as reference in figure 1.

A. System architecture

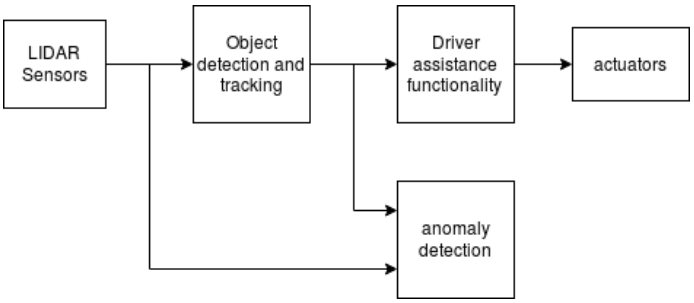

Fig. 1. Components in a LIDAR based driver assistance system

LIDAR sensors supply measurements in form of a point cloud to the *object detection and tracking* component which identifies other traffic participants. Based on this information the *driver assistance functionality* implements the decision making process and controls the actuators. The anomaly detection system aims to detect manipulation of the information that is beeing supplied in form of the *object geometry model*. In the following sections the characteristics (see II-A) and challenges for this anomaly detection problem are detailed.

B. Input Data

Perception algorithms provide driver assistant functions the list of percieved objects further references as *prediction*. An object has a multitude of attributes, of which the most basic ones are:

- position: $\vec{p}$
- velocity: $\vec{v}$
- acceleration: $\vec{a}$
- rotation: $\vec{\theta}$
- dimension: $\vec{dim}$

This information is derived from the *point cloud* generated by the LIDAR sensor, which is represented as a list of points in three dimensional space, further refereced as the set of measurements $M = \{\vec{M}_1, \vec{M}_2, ..., \vec{M}_n\}$. Velocity and acceleration are estimated by accumulating the position of an object incorporating the dimension of time. The input data is characterized as *multivariate* with attributes of *continuous* type and a *spatio-temporal* (see II-A0a) relationship between instances.

C. Type of Anomaly

The type of anomaly we are aiming for to detect are of *collective* and *contextual* nature as a scenario leading to a reaction, fabricated by an attack on the *OGM* (Object Geometry Model), can also occur legitimately and may affect multiple data instances, in which case they are deemed *collective*. For this problem we define the *context* as the *measurement* data supplied by the sensors *M*. If a *prediction P* does not align with the given context (measurement *M*), it is deemed anomalous (see *contextual anomaly* II-A0b).

For this prototype the following attacks will be considered, as they may result in a reaction of the vehicles controlling functions:

1) *Hiding*: Removing an object
2) *Fabrication*: Adding an object
3) *Manipulation*: Changing the positional vector of an object

Reactions that may be provoked include *safety break*, *evasive maneuvers*, *acceleration* and *deceleration*.

D. Data Labels

While it would be possible to automate the creation of labeled data for attack scenarios, such data is not available yet. The automatic manipulation and labeling of data is a high effort task and will need a fair amount of *ground truth* data as well as a formal definition of scenarios. Additionally the task of automation is not trivial as not every scenario can be applied to any ground truth data, and an automatic detection of applicability would be necessary. Therefore we chose an *unsupervised* (see Section II-A0c) approach for this prototype.

E. Output

Because sensorical data is afflicted by *noise* and predictions can only be made with a certain likelihood an anomaly detection mechanism for this problem should assign a *score* representing the probablility of an anomaly, instead of a label (see II-A0d).

IV. Concept

The anomaly detection problem described in III shows the two main issues that need to be addressed:

a) Spatio-temporal data: The relationship between data instances includes the dimension of *time*, as data instances contain *measurements* and *predictions* of the same objects in several different points in *space* and *time*. As the probablility of an anomaly is tied to the specific object which may be manipulated, it is necessary to update the *score* of an object with each incoming data instance, representing a *timestep* (moving forward in the *time* dimension). Therefore we need to model the *time* relationship between data instances, more specifially the impact of past assessments to our current estimation of the probablility of an anomaly. The dimension of *time* is *discrete* for this problem as it is represented by the sequence of incoming data instances which are published with a certain *frequency* which results in discretization of the time component.

b) Contextual anomalies: We defined the *context* for this anomaly detection problem as the raw sensorical data (point cloud, see III-C). The relationship between *contextual* and *behavioral* attributes need to be modeled in order to evaluate the condition under which a data instances is deemed anomalous.

Furthermore there are non functional aspects which should be considered in the solution, depending on the use case of the anomaly detection mechanism. As this problem is allocated in the cyberphysical automotive domain, *safety* is an aspect of huge importance. Therefore the solution must have high *reliability*. Also the *real-time* environment must be considered assuming an *online* use case in which the acting components must include the result of the anomaly detection component in the decision if it should act upon the given situation.

A. Binary Bayes Filter

To handle the *time* relationship between data instances we use a *Binary Bayes Filter* in our prototype. In this section we have a closer look at the mathematical derivation of this filter and how it applies to our problem, with the inputs described in section II-A0a. In the following section the contextual input attribute *measurement* at timestep *t* will be described as m_t, and the behavioural input attribute *prediction* at timestep *t* as o_t. A subset (sequence) of measurements (*M*) is denoted as $m_{m:n}$, the same notation is used for a subset (sequence) of predictions ($o_{m:n}$).

We define *A*, the conditional probability that the prediction o_t of sequence $o_{1:t}$ is anomalous given the contextual input of measurements $m_{1:t}$ as:

$$p(A_t|m_{1:t}, o_{1:t}) \tag{1}$$

Furthermore we assume that what we are going to measure at *t* is not affected by past measurements *1:t-1* (Markov property).

$$p(m_t|A, m_{1:t-1}, o_{1:t}) = p(m_t|A, o_t) \tag{2}$$

And also that the prior probablity of anomaly, independent from measurement at *t*, is also independent from the prediction at *t*, as we need the contextual information (*measurement*) to determine the probablity.

$$p(A|m_{1:t-1}, o_{1:t}) = p(A|m_{1:t-1}, o_{1:t-1}) \tag{3}$$

Using bayes theroem to expand (1) and applying assumptions (2) and (3) this equation is expanded to the following:

$$p(A|o_{1:t}, m_{1:t}) = \frac{p(m_t|A, o_t)p(A|m_{1:t-1}, o_{1:t-1})}{p(m_t|m_{1:t-1}, o_{1:t})} \tag{4}$$

Where $p(m_t|A, o_t)$ represents the *measurement model*, the probability of measuring m_t given an anomalous o_t. Again, expanding with Bayes Theorem and applying the Markov Property:

$$p(A|o_{1:t}, m_{1:t}) = \frac{p(m_t|A, o_t)p(A|m_{1:t-1}, o_{1:t-1})}{p(m_t|m_{1:t-1}, o_{1:t})}$$

$$p(m_t|A, o_t) = \frac{p(A|m_t, o_t)p(m_t|o_t)}{p(A|o_t)}$$

$$p(A|o_{1:t}, m_{1:t}) = \frac{p(A|m_t, o_t)p(m_t|o_t)p(A|m_{1:t-1}, o_{1:t-1})}{p(A|o_t)p(m_t|m_{1:t-1}, o_{1:t})}$$

$$p(A|o_{1:t}, m_{1:t}) = \frac{p(A|m_t, o_t)p(m_t|o_t)p(A|m_{1:t-1}, o_{1:t-1})}{p(A|o_t)p(m_t|m_{1:t-1}, o_{1:t})}$$

results in the following equation:

$$p(A|o_{1:t}, m_{1:t}) = \frac{p(A|m_t, o_t)p(m_t|o_t)p(A|m_{1:t-1}, o_{1:t-1})}{p(A)p(m_t|m_{1:t-1}, o_{1:t})}$$

Because A is a binary state we can formulate

$$p(\neg A|o_{1:t}, m_{1:t}) = \frac{p(\neg A|m_t, o_t)p(m_t|o_t)p(\neg A|m_{1:t-1}, o_{1:t-1})}{p(\neg A)p(m_t|m_{1:t-1}, o_{1:t})}$$

and derive the ratio of probabilities (*Chance*)

$$\frac{p(A|o_{1:t}, m_{1:t})}{p(\neg A|o_{1:t}, m_{1:t})} = \frac{\frac{p(A|m_t, o_t)p(m_t|o_t)p(A|m_{1:t-1}, o_{1:t-1})}{p(A)p(m_t|m_{1:t-1}, o_{1:t})}}{\frac{p(\neg A|m_t, o_t)p(m_t|o_t)p(\neg A|m_{1:t-1}, o_{1:t-1})}{p(\neg A)p(m_t|m_{1:t-1}, o_{1:t})}}$$

$$\frac{p(A|o_{1:t}, m_{1:t})}{p(\neg A|o_{1:t}, m_{1:t})} = \frac{\frac{p(A|m_t, o_t)\cancel{p(m_t|o_t)}p(A|m_{1:t-1}, o_{1:t-1})}{p(A)\cancel{p(m_t|m_{1:t-1}, o_{1:t})}}}{\frac{p(\neg A|m_t, o_t)\cancel{p(m_t|o_t)}p(\neg A|m_{1:t-1}, o_{1:t-1})}{p(\neg A)\cancel{p(m_t|m_{1:t-1}, o_{1:t})}}}$$

$$\frac{p(A|o_{1:t}, m_{1:t})}{p(\neg A|o_{1:t}, m_{1:t})} = \frac{\frac{p(A|m_t, o_t)p(A|m_{1:t-1}, o_{1:t-1})}{p(A)}}{\frac{p(\neg A|m_t, o_t)p(\neg A|m_{1:t-1}, o_{1:t-1})}{p(\neg A)}}$$

$$\frac{p(A|o_{1:t}, m_{1:t})}{p(\neg A|o_{1:t}, m_{1:t})} = \frac{p(A|m_t, o_t)p(A|m_{1:t-1}, o_{1:t-1})p(\neg A)}{p(\neg A|m_t, o_t)p(\neg A|m_{1:t-1}, o_{1:t-1})p(A)}$$

$$\frac{p(A|o_{1:t}, m_{1:t})}{1 - p(A|o_{1:t}, m_{1:t})} = \tag{5}$$

$$\underbrace{\frac{p(A|m_t, o_t)}{1 - p(A|m_t, o_t)}}_{\text{current timestep}} \underbrace{\frac{p(A|m_{1:t-1}, o_{1:t-1})}{1 - p(A|m_{1:t-1}, o_{1:t-1})}}_{\text{recursion}} \underbrace{\frac{1 - p(A)}{p(A)}}_{\text{prior}} \tag{6}$$

to convert the probability to log odds form:

$$l(A|m_{1:t}, o_{1:t}) = \underbrace{l(A|m_t, o_t)}_{\text{model}} + \underbrace{l(A|m_{1:t-1}, o_{1:t-1})}_{\text{recursion}} - \underbrace{l(A)}_{\text{prior}}$$

leading to the final equation of:

$$l_t = model(A, o_t, m_t) + l_{t-1} - l_0 \tag{7}$$

Using log odds eases the process of updating the anomaly score with new evidence. We chose to use log odds over probability as it will help to satisfy possible *real-time* constraints.

B. Jaccard Index

As mentioned in the beginning of this section, the other issue to be solved is the relationship between *contextual* and *behavioural* attributes to evaluate the condition for which a prediction, the behavioural attribute, is deemed to be anomalous. The concept is based on the idea that the sets of measurements and the sets predictions has to have a certain degree of *similarity* for the prediction to be considered legitimate. A statistical method to model similarity

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|} \tag{8}$$

Fig. 2. Jaccard Index

(or diversity) of finite sets is the *Jaccard Index* also called *Intersection over Union*.

To apply this method to our data we need to compute the *intersection* and *union* operations on the sets of *measurements* and *predictions*. Because the elements of these sets are of different types, voxel for the measurements, and bounding boxes for the predictions we need to use a mapping function that lets us compute those operations on these two different types of data, either mapping one to the other or find a common representation. A mapping function for this case translates to the *reverse measurement model* describing what is likely to be measured given a prediction.

Ideally the reverse measurement model would be able to describe the exact voxels that must have been scanned for a given prediction. As this is virtually impossible and also for simplicities sake, we will use a mapping of the two types to a common type, in this protoype a rectangle. This is a very straightforward solution as the *Jaccard Index* can be applied easily to the geometric representation of a rectangle, and we can find an simple function for each type that can map the data to the rectangle. For the the prediction we simlpy remove the height dimension of the bounding box. And for the voxel data we need to fit a minimal rectangle around the measurement points that are probably relevant to the prediction. The relevant voxels are determined by the distance to the center of the predicted object in question, which depends on its dimensions.

V. Evaluation

For the evaluation of the concept we modified the OGM according to the scenarios in III-C.

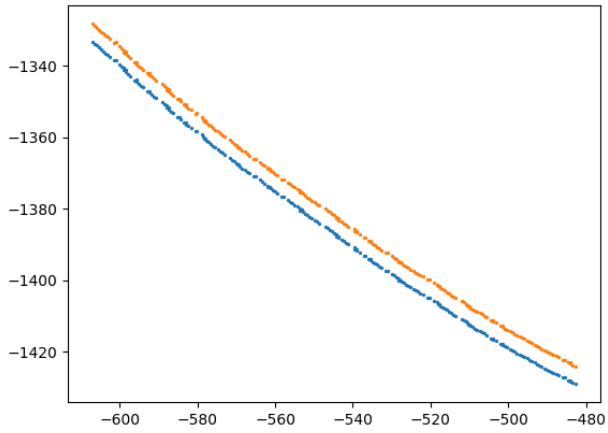

Fig. 3. Car Trajectories

For the *Fabrication* scenario a car has been added, driving with an y-offset of 5 meters, in the coordinate system relative to the ego vehicle, alongside an unmanipulated ground truth object following the same trajectory. In figure 3 the trajectories of the car are displayed using the world coordinate system used by the tracking algorithms. The blue points represent the center point of the ground truth object while the yellow points is the center point of the fabricated car. Using the manipulated data we ran the prototype of the described concept and recorded the anomaly score together with each incoming pair of *predictions* and *measurements*. In 4 the trajectory

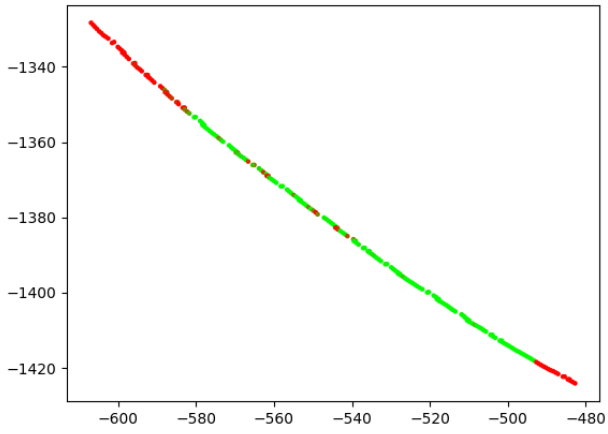

Fig. 4. Anomaly score of the ground truth object

of the unmanipulated car is displayed again, color coded with the anomaly score of each timestep. The coloring is achieved by using RGB coloring in the following way $rgb(1 - anomaly\_score, anomaly\_score, 0)$ while *anomaly_score* is a floating point number ranging from 0 (low confidence) to 1 (high confidence), a low confidence meaning a high probability of anomaly (red). In the scenario the ground truth object is overtaken by the ego vehicle running the tracking algorithms, gradually reducing the distance between the two vehicles. When performing measurements at high distance scan points are most dense at the rear end of the car and a scan beam is unlikely to hit the A-pillar of the tracked object. As a result the rectangle that is fit around the scan points in the mapping function will only be covering the rear end of the car resulting in a low intersection over union, and as a result a low confidence, as can be seen in figure 4. As soon as the tracked vehicle is close enough that points in the frontal section of the car are scanned, the fitted rectangle increasingly matches the bounding box of the object and the confidence increases. After passing the car the confidence drops again as the sensors are aligned to measure in front of the vehicle, and there are no more incoming measurements that hit the area of the tracked object.

Figure 5 shows the trajectory and anomaly score of the fabricated object in the same manner.

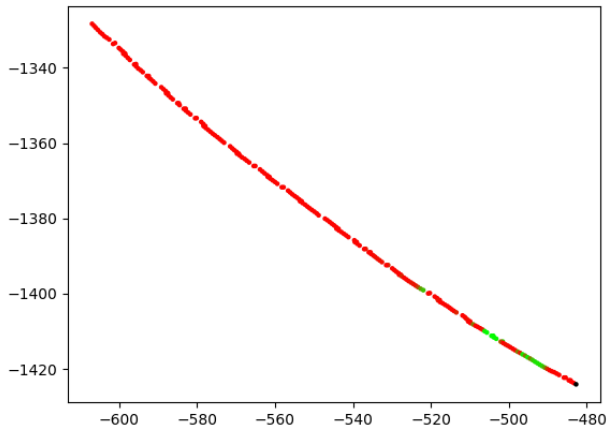

Fig. 5. Anomaly score of the fabricated object

The fabricated object is successfully detected as an anomaly until the distance to the overtaking is only a few meters. This is due to scan points hitting the ground in front of the car performing the tracking. These points are erroneously associated with the fabricated object leading to a rise in confidence.

Another scenario that has been evaluation is the manipulation of an objects position as shown in 6

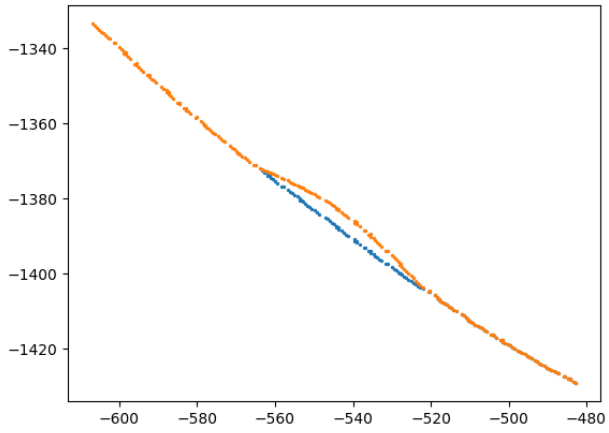

Fig. 6. Car Trajectories

Here the trajectory of an existing ground truth object has been manipulated for a certain period. Blue points indicate the original real trajectory. In figure 7 the outcome of the anomaly detection is displayed again showing a low confidence for the manipulated trajectory, but also showing the same issues related to the very basic mapping function used in this prototype as discussed before.

The last scenario that was mentioned in III-C is *Hiding*, deletion of an object. The mapping function used in this

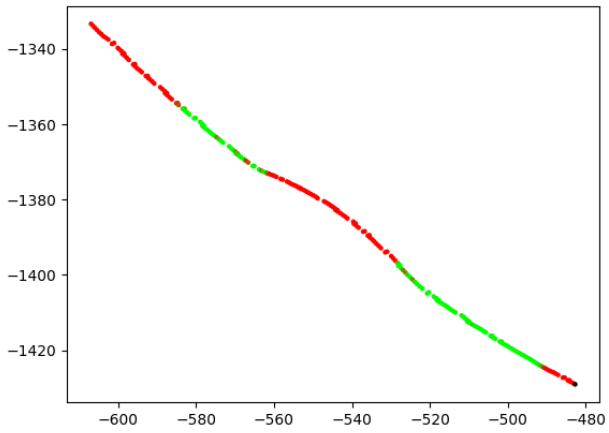

Fig. 7. Anomaly Score.

prototype has the constraint that it needs the objects location and dimension to map the measurement to the common representation of a rectangle. Therefore it is not possible to detect the deletion of an object with this approach.

VI. Conclusion and Outlook

While the basic concept has proven to be successfull the realization needs improvement to be reliable and to satisfy every functional requirement. For this prototype a very straightforward and easy solution circumventing the need for a real *reverse measurement model* was chosen, by mapping the measurement and the prediction to a rectangle, reducing the dimensions of the data drastically. When using a *reverse measurement model* that is derived from the *sensor model* used to create the measurements results will improve and issues described the evaluation can be negated.

While this approach is based on the concept of confirming occupied space with measurements it is not sufficient to detect deleted objects. For this functional requirement *free space detection* can be incoorporated in the solution. Also the Binary Bayes filter used in this prototype is derived for usage of a single source of *evidence*. This can be adapted to use multiple sources of evidence as contextual attributes.

A car is usually equipped with more than one sensor, while data from multiple LIDAR sensors are used for this prototype the information of origin is discarded, again reducing the dimensions of input data. Modeling relationships can help identify anomalies by incooporating this information into the anomaly detection mechanism.

References

[1] J. Kittler, W. Christmas, T. de Campos, D. Windridge, F. Yan, J. Illingworth, and M. Osman, "Domain anomaly detection in machine perception: A system architecture and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 5, pp. 845–859, May 2014.

[2] A. Patcha and J.-M. Park, "An overview of anomaly detection techniques: Existing solutions and latest technological trends," *Computer Networks*, vol. 51, no. 12, pp. 3448 – 3470, 2007. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S138912860700062X

[3] S. Thrun and A. Bücken, "Integrating grid-based and topological maps for mobile robot navigation," in *Proceedings of the AAAI Thirteenth National Conference on Artificial Intelligence*, Portland, Oregon, 1996.

[4] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, 07 2009.

[5] T. M. Hospedales, J. Li, S. Gong, and T. Xiang, "Identifying rare and subtle behaviors: A weakly supervised joint topic model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 12, pp. 2451–2464, Dec 2011.

[6] D. Weinshall, A. Zweig, H. Hermansky, S. Kombrink, F. Ohl, J. Anemüller, J.-H. Bach, L. Van Gool, F. Nater, T. Pajdla, M. Havlena, and M. Pavel, "Beyond novelty detection: Incongruent events, when general and specific classifiers disagree," *IEEE transactions on pattern analysis and machine intelligence*, vol. 34, 12 2011.

[7] A. Zweig and D. Weinshall, "Exploiting object hierarchy: Combining models from different category levels," in *2007 IEEE 11th International Conference on Computer Vision*, Oct 2007, pp. 1–8.