Deep-Learning-Based Anomaly Detection for Lane-Changing Decisions

Sheng-Li Wang[1], Chien Lin[1], Srivalli Boddupalli[2], Chung-Wei Lin[1], Sandip Ray[2]
[1]National Taiwan University, [2]University of Florida
{r09922039, r10922a06}@ntu.edu.tw, bodsrivalli12@ufl.edu, cwlin@csie.ntu.edu.tw, sandip@ece.ufl.edu

***Abstract*— Vehicles can utilize their sensors or receive messages from other vehicles to acquire information about the surrounding environments. However, the information may be inaccurate, faulty, or maliciously compromised due to sensor failures, communication faults, or security attacks. The goal of this work is to detect if a lane-changing decision and the sensed or received information are anomalous. We develop three anomaly detection approaches based on deep learning: a classifier approach, a predictor approach, and a hybrid approach combining the classifier and the predictor. All of them do not need anomalous data nor lateral features so that they can generally consider lane-changing decisions before the vehicles start moving along the lateral axis. They achieve at least 82% and up to 93% F_1 scores against anomaly on data from Simulation of Urban MObility (SUMO) [1] and HighD [2]. We also examine system properties and verify that the detected anomaly includes more dangerous scenarios.**

I. Introduction

Lane changing is a common maneuver on roads, but it is challenging for drivers. Many Advanced Driver Assistance Systems (ADAS) are developed to assist lane-changing maneuvers nowadays. Vehicles can utilize their sensors or receive messages from other vehicles to acquire information about the surrounding environments. However, the information which leads the vehicles to make decisions and perform maneuvers may be inaccurate, faulty, or maliciously compromised due to sensor failures, communication faults, or security attacks. As vehicles are safety-critical systems, it is crucial to detect the inaccurate, faulty, or compromised information and identify the corresponding vehicles which behave anomalously to improve the overall system safety and robustness.

Related work includes lane-changing detection and anomaly detection. For the lane-changing detection problem, existing work can be categorized into three categories: rule-based approaches [3], [4], [5], [6], probabilistic models largely with hidden Markov models [7], [8], [9], and deep learning approaches based on recurrent neural networks [10], [11], [12], [13], [14], [15] and other neural networks [16], [17], [18]. For the anomaly detection problem, learning the distribution of normal data is the key concept [19], [20]. Section II will provide a detailed review.

The perspective of the related work is vehicle-to-vehicle: through detecting or predicting other vehicles' maneuvers or intentions, the ego vehicle can make its decision. This work takes a different and more general perspective on the anomaly detection problem. The goal of this work is to *detect if a lane-changing decision and the sensed or received information are anomalous* (or inconsistent). If they are anomalous, it implies that either the lane-changing decision is unreasonable or there is anomalous information, no matter the information is sensed from sensors or received from communication, and no matter the anomaly detection is performed on a vehicle or an outsider, *e.g.*, roadside unit or infrastructure. As a result, the detection approaches can be applied to vehicles or different outsiders, and they can take the corresponding reactions accordingly, *e.g.*, runtime decision or offline maintenance.

In this work, our main contributions include

- We develop three anomaly detection approaches based on deep learning: a classifier approach, a predictor approach, and a hybrid approach combining the classifier and the predictor. All of them do not need anomalous data nor lateral features so that they can generally consider lane-changing decisions before the vehicles start moving along the lateral axis.
- We evaluate our approaches with Simulation of Urban MObility (SUMO) [1] and HighD [2]. The proposed approaches achieve at least 82% and up to 93% F_1 scores against anomaly which is *stealthy* to the laws of physics. Especially, the hybrid approach has the best detection performance.
- We also examine system properties and verify that the detected anomaly includes more dangerous scenarios, and the mis-detected anomaly includes less dangerous scenarios. This is crucial as a fundamental goal of system design is to prevent collisions, and a detection approach should focus more on those dangerous scenarios.

The rest of this paper is organized as follows. Section II reviews related work. Section III describes our problem formulation. Section IV presents the proposed approaches. Section V demonstrates and discusses the experimental results. Section VI concludes this paper.

II. Related Work

The lane-changing detection problem has been widely studied in the past decades. The approaches can be distinguished by three categories: rule-based approaches, probabilistic models, and deep learning approaches. We introduce them as well as the corresponding anomaly detection in this section.

A. Rule-Based Approaches

Rule-based approaches detect a lane-changing maneuver by setting predefined rules, *e.g.*, laws of physics. They are usually simple, intuitive, and lightweight. A lane-changing maneuver can be detected when a vehicle's lateral speed or acceleration is higher than a threshold [3]. A vehicle's distance from the left or right lane boundary [4] and its steering angle [5] can also be used to detect a lane-changing maneuver. Khelfa and Tordeux introduced a rule-based model, where the speed and position differences of the four surrounding vehicles on the current and target lanes are the inputs of the model [6].

B. Probabilistic Models

A probabilistic model is a statistical technique using past data to predict the probability of the occurrence of a future event. The detection performance depends on chosen features and hidden states of the probabilistic model [21]. Park *et al*. proposed a hidden Markov model to detect a lane-changing maneuver [7]. Li *et al*. further integrated a hidden Markov model with a Gaussian mixture model to form a hybrid model and detect a lane-changing maneuver [8]. Sharma *et al*. combined a continuous hidden Markov model and a discrete hidden Markov model to detect lane-changing maneuvers [9]. The continuous hidden Markov model generates hidden state sequences from trajectory inputs, and the discrete hidden Markov model classifies a sequence into different driving maneuvers. As for fuzzy logic based hidden Markov model, an approach discriminates driving maneuvers into very safe, safe, and dangerous driving scenarios [22]. A Bayes model was also proposed to predict whether a driver will take over the leading car [23]. Ma *et al*. proposed a logistic regression model with different levels of input features to detect the lane-changing intention [24].

C. Deep Learning Approaches

With a huge amount of driving data, deep learning models have become popular to solve the problem. Learning from data leads to a more comprehensive detection capability. We can not only predict whether a lane-changing maneuver will happen but also predict the corresponding position and timing. The problems can be categorized into three categories: the classification problem, the trajectory and position prediction problem, and the hybrid problem [25].

Many existing deep learning approaches applied Recurrent Neural Networks (RNNs) [10], [11], [12], [13], [14], [15]. lane-changing maneuvers are usually based on a few seconds of environmental information, so it is straightforward to use RNN-based models which are well-suited for time series data. Zyner *et al*. used a multi-layer Long Short-Term Memory (LSTM)-based RNN to predict a driver's intention before a road intersection [10]. Park *et al*. proposed a LSTM-based autoencoder to predict a vehicle's trajectory [11]. Xin *et al*. adopted dual LSTM-based RNNs, where one LSTM is for recognizing a driver's intention, and the other one is for predicting a vehicle's trajectory [12]. Yan *et al*. proposed a variant of RNN like a network of gated recurrent units to predict when a driver will shift the lane and when the driver will complete the lane changing [13]. Xing *et al*. predicted a lane-changing intention by an ensemble bi-directional RNN

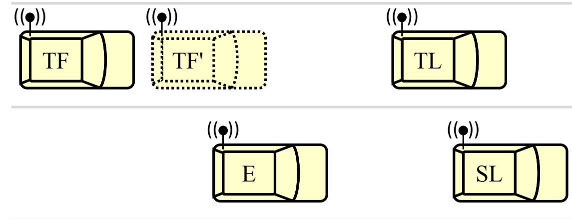

Fig. 1. The lane-changing scenario, where E is the ego vehicle, SL is the leading vehicle on the source lane, TL is the leading vehicle on the target lane, and TF is the following vehicle of on the target lane. Anomalous information about TF may indicate TF as TF′. If E changes its lane, the changing decision should be detected as an unreasonable decision, and thus there is anomalous information.

model with LSTM units [14]. Wirthmüller *et al*. used a single layer LSTM-based RNN to predict the time until a vehicle changes the lane [15] .

There are other deep learning approaches applying other neural networks. The spatial information about vehicle relations can be well captured by convolutional neural networks [26]. Hu *et al*. predicted the intention and the motion of a vehicle under various driving scenarios by a multi-layer fully-connected deep neural network [16]. Xie *et al*. detected a lane-changing process by a deep belief network [17]. De Candido *et al*. used three 1-D convolution-based deep autoencoders on driving maneuvers to detect a driver's maneuver [18].

D. Anomaly Detection

Deep learning approaches are broadly used for anomaly detection [19], [20]. Ramyar *et al*. used a one class support vector machine to detect anomaly for lane-changing maneuvers [27]. Fan *et al*. used a recurrent convolution autoencoder and a one class support vector machine to detect anomaly [28]. Guo *et al*. also used an autoencoder to report anomaly if the reconstruction error is above a certain threshold [29]. We can observe that most anomaly detection approaches used some methods to learn the distribution from normal training data and report anomaly if test data have very different distributions from the trained distributions.

III. Problem Formulation

A. System Model and Anomaly Model

We consider a scenario that an ego vehicle changes its lane, where the ego vehicle senses or receives information and acts accordingly. In this work, we focus on whether there is anomalous information, no matter the information is sensed from sensors or received from communication. This means that anomaly can represent an anomalous maneuver, a wrongly-sensed normal maneuver, or wrongly-received information about a normal maneuver. As a result, the scenario is quite general, not limited to a non-connected (sensors only), partially-connected (sensors and communication), or fully-connected environment (maybe communication only). The scenario is illustrated in Figure 1.

To formally define anomaly, we first define the sensed or received information at time step t as a *feature vector* with dimension n:

$$\mathbf{r}^{(t)} = \left[r_1^{(t)}, r_2^{(t)}, \dots, r_n^{(t)}\right], \tag{1}$$

where a feature can be the location, speed, or acceleration of the ego vehicle, the leading vehicle on the source lane, the leading vehicle on the target lane, or the following vehicle on the target lane. Note that the four vehicles play the most important roles for the lane-changing maneuver of the ego vehicles.

If a lane-changing maneuver happened at time step t, we combine the feature vectors before t with a window size w and form a time series of feature vectors, called a *trajectory vector*:

$$\mathbf{R} = \left[\mathbf{r}^{(0)}, \mathbf{r}^{(1)}, \dots, \mathbf{r}^{(w-1)}\right]. \tag{2}$$

Here, we assume that the acceleration features in $\mathbf{R}$ are:

$$\mathbf{A} = \left[\mathbf{a}^{(0)}, \mathbf{a}^{(1)}, \dots, \mathbf{a}^{(w-1)}\right]. \tag{3}$$

The anomaly adds an offset vector $\mathbf{o}$ to $\mathbf{A}$, where each element in $\mathbf{o}$ is zero or a constant o (the severity of anomaly), and at least one element in $\mathbf{o}$ is o, meaning that the acceleration features of at least one vehicle are changed. As a result. we can get:

$$\begin{aligned}\mathbf{A}' &= \left[\mathbf{a}^{(0)} + \mathbf{o}, \mathbf{a}^{(1)} + \mathbf{o}, \dots, \mathbf{a}^{(w-1)} + \mathbf{o}\right] \\ &= \left[\mathbf{a}'^{(0)}, \mathbf{a}'^{(1)}, \dots, \mathbf{a}'^{(w-1)}\right].\end{aligned} \tag{4}$$

Given $\mathbf{A}'$, we can use the laws of physics to calculate the corresponding speed and location features starting from $\mathbf{a}'^{(0)}$ to $\mathbf{a}'^{(w-1)}$:

$$\mathbf{v}'^{(t+1)} = \mathbf{v}^{(t)} + \mathbf{a}'^{(t)} \cdot \Delta t, \tag{5}$$

$$\mathbf{l}'^{(t+1)} = \mathbf{l}^{(t)} + \mathbf{v}'^{(t)} \cdot \Delta t + \frac{1}{2} \cdot \mathbf{a}'^{(t)} \cdot \Delta t^2, \tag{6}$$

where $\Delta t = 1$ (a time step). Then, by assembling the features, we get the updated anomalous trajectory vector:

$$\mathbf{R}' = \left[\mathbf{r}'^{(0)}, \mathbf{r}'^{(1)}, \dots, \mathbf{r}'^{(w-1)}\right]. \tag{7}$$

By the setting, the anomaly is *stealthy* against the laws of physics as the locations, speeds, accelerations will satisfy (not be detected by) the laws of physics.

For example, if the anomaly adds a negative o to the acceleration of the leading vehicle on the target lane and computes the corresponding $\mathbf{R}'$, implying that the vehicle is slower. In this case, if the ego vehicle still changes its lane, the changing decision should be detected as an unreasonable decision, and thus there is anomalous information.

We will also examine system properties with the anomaly to verify if the anomaly is worth being detected, *i.e.*, if the anomaly has no or very little effects on system properties, then it is not worth being detected.

B. Detection Goal

Given a lane-changing trajectory vector $\mathbf{R}$, an anomaly detection approach can be represented as a function:

$$F(\mathbf{R}) = \begin{cases} 0, & \text{there is no anomaly in } \mathbf{R}; \\ 1, & \text{there is anomaly in } \mathbf{R}. \end{cases} \tag{8}$$

In other words, if $F(\mathbf{R}) = 0$, the anomaly detection approach considers that a lane-changing decision with $\mathbf{R}$ is reasonable; if $F(\mathbf{R}) = 1$, the anomaly detection approach considers that a lane-changing decision with $\mathbf{R}$ is unreasonable, based on anomalous information.

We use the F_1 score to evaluate the detection performance. It is computed as follows:

$$\text{TP} = |\{\mathbf{R}' \mid F(\mathbf{R}') = 1\}|, \tag{9}$$

$$\text{FN} = |\{\mathbf{R}' \mid F(\mathbf{R}') = 0\}|, \tag{10}$$

$$\text{FP} = |\{\mathbf{R} \mid F(\mathbf{R}) = 1\}|, \tag{11}$$

$$\text{TN} = |\{\mathbf{R} \mid F(\mathbf{R}) = 0\}|, \tag{12}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \tag{13}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \tag{14}$$

$$F_\beta \text{ Score} = (1 + \beta^2) \cdot \frac{\text{Precision} \cdot \text{Recall}}{(\beta^2 \cdot \text{Precision}) + \text{Recall}}. \tag{15}$$

In other words, TN (true negative) is the number of normal trajectory vectors which are correctly classified as normal trajectory vectors; FP (false positive) is the number of normal trajectory vectors which are wrongly detected as anomalous trajectory vectors; TP (true positive) is the number of anomalous trajectory vectors which are correctly detected as anomalous trajectory vectors; FN (false negative) is the number of anomalous trajectory vectors which are wrongly classified as normal trajectory vectors. We set $\beta = 1$, in which precision is considered the same important as recall.

It should be mentioned here that a feature vector or a trajectory vector does not need lateral features such as the distance to the line between the source and target lanes. If there is no lateral feature, an anomalous trajectory vector describes an unreasonable lane-changing decision, considering the longitudinal relations between vehicles. This is different from the related work as we can consider the potential lane-changing decisions before the vehicles start moving along the lateral axis. Nevertheless, if there are lateral features, the formulation and the proposed approaches are still applicable.

IV. Proposed Approaches

In this section, we develop four approaches to perform anomaly detection for lane-changing decisions. They are a rule-based approach (mainly for comparison), a classifier approach, a predictor approach, and a hybrid approach. The last three approaches are all learning-based, and they are not limited to any specific attack since we do not use anomalous data during training.

A. Rule-Based Approach

A rule-based approach is based on some laws of physics or some observations of system properties. It is usually difficult to cover all possible scenarios, as some lane-changing maneuvers are complicated such as tactical lane changing or cooperative lane changing [30]. Here, we propose a rule-based approach mainly for comparison. We observe that a reasonable lane-changing maneuver usually has the following features:

- It has no sudden brake after changing lane.
- It does not force the following vehicle on the target lane to have a sudden brake.

For safety reasons, the speed differences between the ego vehicle and its leading and following vehicles on the target lane is better to be close to 0. Otherwise, changing the speed during the lane-changing maneuver may confuse other vehicles and reduce safety or efficiency. Therefore, we design a rule-based detection approach which checks the speed differences. If the speed difference between the ego vehicle and its leading or following vehicles on the target lane is larger than a threshold, it is identified as anomaly.

B. Classifier Approach

The workflow of the classifier approach is illustrated in Figure 2. Here, we treat the anomaly detection as a binary classification problem. During training, we use normal data (no anomalous data) and label lane keeping (no lane changing) as 0 and lane changing as 1. This means that, if there is a lane-changing maneuver at time step t, we have lane-changing data (with the label 1):

$$\mathbf{R_c} = \left[\mathbf{r}^{(t-w)}, \mathbf{r}^{(t-w+1)}, \ldots, \mathbf{r}^{(t-1)}\right]. \tag{16}$$

With a stride s, we have lane-keeping data (with the label 0):

$$\mathbf{R_k} = \left[\mathbf{r}^{(t-s-w)}, \mathbf{r}^{(t-s-w+1)}, \ldots, \mathbf{r}^{(t-s-1)}\right]. \tag{17}$$

Note that the index of time can be shifted to match Equation (2) in Section III.

In the classification, the two classes are both meaningful. The lane-keeping class not only represents lane-keeping maneuvers but also indicates that the surrounding environments are not suitable for lane changing. This is because a vehicle does not suddenly change its lane as it first observes the surrounding environment before making a decision. It chooses to change its lane at a specific time step very likely because the previous time step is not suitable for lane changing. During testing, if we correctly classify lane changing or lane keeping, the corresponding data are regarded as normal data (no anomaly); otherwise, they are regarded as anomaly.

We propose two model structures to develop our classifier approach. The first one is a fully connected Deep Neural Network (DNN) model, and the second one is a Long Short-Term Memory (LSTM) model. In the DNN model, we flatten the time series of trajectory vectors $\mathbf{R}$ as the input. In the LSTM model, we directly take the time series data $\mathbf{R}$ as the input. Both models are binary classifiers and output 0 or 1 representing their classification results.

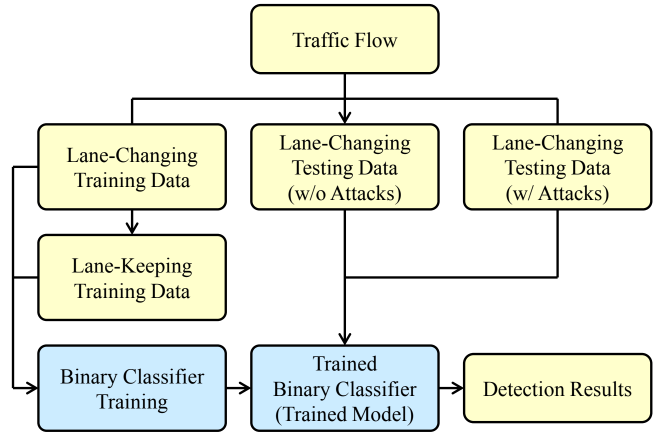

Fig. 2. The workflow of the classifier approach. The binary classifier classifies lane-changing and lane-keeping maneuvers.

C. Predictor Approach

We use normal lane-changing data and predict the trajectory after a vehicle changes its lane. Given an input trajectory vector as Equation (16), we use a predictor to predict the trajectory:

$$\mathbf{R_p} = \left[\mathbf{r}^{(t+1)}, \mathbf{r}^{(t+2)}, \ldots, \mathbf{r}^{(t+w)}\right]. \tag{18}$$

after the lane-changing maneuver at time step t, where w is the window size. If the prediction error (distance from the sensed or received trajectory vector) is larger than a threshold, the corresponding data are regarded as anomaly. We also use a LSTM model as our predictor since the trajectories are more time dependant, and thus it is expected to perform better than a simple DNN model.

D. Hybrid Approach

One benefit of the classifier approach is that it provides good interpretation of the problem, compared with the predictor approach, but it also has some limitations due to some special cases in training data. For example, if a vehicle reacts slowly, its lane-keeping data may still imply a suitable moment for lane changing. On the other hand, if a vehicle acts dangerously, its lane-changing data may still mean a non-suitable moment for lane changing. These cases confuse the classifier and lead to the reduction of detection performance. Due to the safety nature of automotive systems, undetected anomaly maneuvers (false negatives) are more critical. Therefore, we propose a hybrid approach utilizing this assumption to further improve the detection performance. The hybrid approach is a two-phase approach which combines the classifier approach and the predictor approach. In the first phase, we use the classifier approach, and there are two possible outcomes: negative or positive. If it is a negative outcome, we use the predictor approach in the second phase; otherwise (a positive outcome), the outcome is kept to remain the same and considered as anomaly. In other words, the hybrid approach keep positive outcomes and probably turn some negative outcomes into positive outcomes.

V. Experiment Results

A. Experimental Setup

We evaluate our approaches with Simulation of Urban MObility (SUMO) [1] and HighD [2] as our datasets. SUMO is a simulation platform where we can simulate connected environments and collect simulation traces as a dataset. HighD is a dataset collected on German highways by drones, consisting of vehicle trajectories controlled by human drivers.

We use some normal lane-changing data as the training data. We use the method described in Section III to create anomalous data and combine other normal data and anomalous data to form the testing data. For SUMO, we build one 1,000-meter straight road with three lanes to simulate highway scenarios. To observe lane-changing maneuvers more efficiently, some vehicles are set as on-ramp vehicles which try to change to inner lanes, and some vehicles are set as off-ramp vehicles which try to change to outer lanes. The training data and testing data are collected with 800,000 simulation steps and 10,000 simulation steps, respectively. There are 86,655 lane-changing maneuvers for training and 1,214 lane-changing maneuvers for testing. For HighD, the data are collected from six locations, and we take 10% of the lane-changing maneuvers as the normal testing data.

We compare the rule-based approach (RBS), the classifier approach (DCLF or LCLF), the predictor approach (PDT), and the hybrid approach (DHBD or LHBD), where “D” or “L” with CLF and HBD means that we use the DNN or LSTM model, respectively, for the classifier approach. There detailed settings are as follows:

- For the classifier model trained with SUMO, the input window size is 20, and the simulation step is 0.1 second, which means that the model uses 2-second information to make a lane-changing decision.
- For the predictor model trained with SUMO, the input window size is 20, the output window size is 7, and the simulation step is 0.1 second, which means that the model uses 2-second information to decide the later 0.7-second trajectory.
- For the classifier model trained with HighD, the input window size is 75 frames, the number of input frames is 25 per second, which means that the model uses 3-second information to make a lane-changing decision.
- The predictor model trained with HighD uses 3-second information to decide the later 0.7-second trajectory. Benterki *et al*. mentioned that important features change after 3 seconds on average [31], which is consistent within our experiments.
- For SUMO, the stride is 10 simulation steps (1 second) in Equation (17). For HighD, the stride is 50 frames (2 seconds) in Equation (17). Note that, since HighD is a real-world dataset with human drivers, the decision is assumed to be slower.

The neural network structures are implemented using the PyTorch library. All the experiment runs on a desktop with Intel Core i7-9700 CPU and NVIDIA-2080Ti GPU.

TABLE I
EXPERIMENTAL RESULTS (F_1 SCORES) WITH SUMO.

Anomaly	RBS	DCLF	LCLF	PDT	DHBD	LHBD
Leading	0.75	0.77	0.80	0.84	0.86	0.88
Following	0.78	0.82	0.80	0.84	0.85	0.84
Both	0.86	0.92	0.92	0.92	0.93	0.93

TABLE II
EXPERIMENTAL RESULTS (F_1 SCORES) WITH HIGHD.

Anomaly	RBS	DCLF	LCLF	PDT	DHBD	LHBD
Leading	0.77	0.76	0.77	0.78	0.83	0.83
Following	0.73	0.75	0.81	0.90	0.89	0.91
Both	0.83	0.82	0.86	0.90	0.92	0.93

B. Experimental Results with SUMO

We create anomaly on the leading vehicle on the target lane, the following vehicle on the target lane, and both. The experimental results are listed in Table I. The rule-based approach actually achieves good detection performance, but the rules need to be redefined for different driving scenarios. The classifier approach is slightly better than the rule-based approach. In SUMO, the boundary between lane changing and lane keeping is clear since it is a controlled simulation system. However, the decision-making features are not limited to our features. For example, sometimes a vehicle makes a decision which is the best for the overall environment, not itself [30]. Due to this reason, the lane-keeping data do not necessarily mean a non-suitable moment for lane changing. The predictor approach has even better detection performance than the rule-based approach and the classifier approach. This is probably because it tries to identify which is the actual leading vehicle to make the prediction error small. The predictor approach implicitly does similar things as the classifier approach, but it considers more with trajectories and leads to a better detection performance. The hybrid approach further improves the detection performance. The hybrid approach is similar to an ensemble method which gives a more robust detection performance. Another observation is that the DNN model and the LSTM model have similar detection performance, and the LSTM model is slightly better as it can better process time series data.

C. Experimental Results with HighD

We first take a look at the detection performance of some previous work. Two existing references [6] and [24] also tested on HighD and used similar feature spaces as ours which does not include lateral features. Each of them achieves 80% classification accuracy $\left(\frac{\text{TP+TN}}{\text{TP+TN+FP+FN}}\right)$ for lane-changing maneuvers. We believe that it is difficult to achieve 99% accuracy with HighD as the reference [18] if there is no lateral feature. However, as mentioned in Section III, we are targeting anomalous or unreasonable lane-changing decisions, considering the longitudinal relations between vehicles, before the vehicles start moving along the lateral axis.

Similarly, we create anomaly on the leading vehicle on the target lane, the following vehicle on the target lane, and both.

The experimental results are listed in Table II. Similar to the experimental results with SUMO, the predictor approach and the hybrid approach outperform the rule-based approach and the classifier approach, showing that the results are not limited to simulation data. Compared with the experimental results with SUMO, when the anomaly is on the leading vehicle of the target lane, the detection performance of most approaches decrease a little, but the trend is not clear when the anomaly is on the following vehicle of the target lane.

The classifier approach does not have significant improvement over the rule-based approach. This is because HighD is a real-world dataset with human drivers, and human drivers are not as rational as the controllers in SUMO. Similar to the results with SUMO, the predictor approach has better detection performance than the rule-based approach and the classifier approach. Although the classifier approach does not have the same performance as the predictor approach, it has an advantage that it can output detection results earlier than the predictor approach. In addition, they can be combined as the hybrid approach which further improves the detection performance. On the other hand, it can also be observed that the LSTM model has better detection performance, compared with the DNN model.

D. System Properties

We also examine some system properties with the anomaly to verify if the anomaly is worth being detected, *i.e.*, if the anomaly has no or very little effects on system properties, then it is not worth being detected. Liu *et al.* evaluated lane-changing maneuvers based on four metrics [32]: the success rate of lane changing, the minimum distance gap, the total fuel cost which is the integral of the square of acceleration, and the accumulated discomfort which is the integral of the square of jerk. In this work, we focus on lane-changing trajectory vectors, so there is no need to discuss the success rate of lane changing. Also, the fuel cost and the discomfort highly depend on the acceleration as well as the speed difference. Therefore, we evaluate the following system properties:

- The mean gaps to the leading vehicle and the following vehicle on the target lane after the lane-changing maneuver.
- The mean speed differences from the leading vehicle and the following vehicle on the target lane after the lane-changing maneuver.
- The Time To Collision (TTC) which is the time until a collision between two vehicles if the courses and the speed difference are maintained [33]. A smaller TTC implies a more dangerous scenario.

Table III lists the system properties with SUMO, anomaly on the following vehicle on the target lane, and the hybrid approach with the LSTM model. There are some observations as follows, and similar observations can also be found with different attacked vehicle as well as with HighD. First, as the anomaly is on the following vehicle on the target lane, the system properties related to the leading vehicle on the target lane do not have significant differences. Second, the TP (true positive) scenarios have the smallest mean gap to the following vehicle, while the FN (false negative) scenarios have much larger mean gap, compared with the TP (true positive) scenarios. This means that the detected anomaly indeed includes more dangerous scenarios, and the mis-detected anomaly includes less dangerous scenarios. Third, the mean speed difference has a similar trend. The detected anomaly includes more dangerous scenarios, where the following vehicle on the target lane needs to perform hard breaks to prevent collisions, while the mis-detected anomaly includes less dangerous scenarios. Last, the detected anomaly includes 15% scenarios where the TTC is less than 2 second. Again, this demonstrates that the detected anomaly includes more dangerous scenarios, where the following vehicle on the target lane needs to perform hard breaks to prevent collisions.

TABLE III
SYSTEM PROPERTIES WITH SUMO, ANOMALY ON THE FOLLOWING VEHICLE ON THE TARGET LANE, AND THE LHBD APPROACH.

System Property	TP	FN	FP	TN
Gap: Leading (m)	101.35	94.84	108.00	97.56
Gap: Following (m)	49.61	117.49	145.63	127.86
Diff. Speed: Leading (m/s)	3.19	3.78	3.16	3.28
Diff. Speed: Following (m/s)	-4.35	-2.61	-0.17	-0.61
Ratio of [TTC $<$ 2s]	15%	0%	0%	0%

VI. CONCLUSIONS

In this work, we developed three anomaly detection approaches based on deep learning: a classifier approach, a predictor approach, and a hybrid approach. All of them do not need anomalous data nor lateral features so that they can generally consider lane-changing decisions before the vehicles start moving along the lateral axis. They achieved at least 82% and up to 93% F_1 scores against anomaly, which is stealthy to the laws of physics, on data from SUMO [1] and HighD [2]. We also examined system properties and verified that the detected anomaly includes more dangerous scenarios. Future directions include more complicated lane-changing scenarios, detection based on other learning models, and reaction after detecting anomaly.

ACKNOWLEDGEMENT

This work is partially supported by National Science Foundation (NSF) under Grant Number CNS-1908549, Ministry of Education (MOE) in Taiwan under Grant Numbers NTU-109V0901 and NTU-110V0901, Ministry of Science and Technology (MOST) in Taiwan under Grant Numbers MOST-110-2636-E-002-026 and MOST 111-2636-E-002-018, and Qualcomm.

REFERENCES

[1] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using SUMO," in *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 2575–2582.

[2] R. Krajewski, J. Bock, L. Kloeker, and L. Eckstein, "The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems," in *International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 2118–2125.

[3] J. Nilsson, J. Fredriksson, and E. Coelingh, "Rule-based highway maneuver intention recognition," in *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2015, pp. 950–955.

[4] J. Lioris, A. Bracquemond, G. Thiolon, and L. Bonic, "Lane change detection algorithm on real world driving for arbitrary road infrastructures," in *IEEE Annual Computer Software and Applications Conference (COMPSAC)*, vol. 1, 2018, pp. 561–569.

[5] P. Leakkaw and S. Panichpapiboon, "Real-time lane change detection through steering wheel rotation," in *IEEE Vehicular Networking Conference (VNC)*, 2018, pp. 1–7.

[6] B. Khelfa and A. Tordeux, "Lane-changing prediction in highway: Comparing empirically rule-based model MOBIL and a naïve Bayes algorithm," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021, pp. 1598–1603.

[7] S. Park, W. Lim, and M. Sunwoo, "Robust lane-change recognition based on an adaptive hidden Markov model using measurement uncertainty," *International Journal of Automotive Technology*, vol. 20, pp. 255–263, 2019.

[8] J. Li, N. Gattu, and S. Ghosh, "FAuto: An efficient GMM-HMM FPGA implementation for behavior estimation in autonomous systems," in *International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–8.

[9] O. Sharma, N. C. Sahoo, and N. B. Puhan, "Highway discretionary lane changing behavior recognition using continuous and discrete hidden markov model," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021, pp. 1476–1481.

[10] A. Zyner, S. Worrall, J. Ward, and E. Nebot, "Long short term memory for driver intent prediction," in *IEEE Intelligent Vehicles Symposium (IV)*, 2017, pp. 1484–1489.

[11] S. H. Park, B. Kim, C. M. Kang, C. C. Chung, and J. W. Choi, "Sequence-to-sequence prediction of vehicle trajectory via LSTM encoder-decoder architecture," in *IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 1672–1678.

[12] L. Xin, P. Wang, C.-Y. Chan, J. Chen, S. E. Li, and B. Cheng, "Intention-aware long horizon trajectory prediction of surrounding vehicles using dual lstm networks," in *International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 1441–1446.

[13] Z. Yan, K. Yang, Z. Wang, B. Yang, T. Kaizuka, and K. Nakano, "Time to lane change and completion prediction based on gated recurrent unit network," in *IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 102–107.

[14] Y. Xing, C. Lv, H. Wang, D. Cao, and E. Velenis, "An ensemble deep learning approach for driver lane change intention inference," *Transportation Research Part C Emerging Technologies*, vol. 115, p. 102615, 2020.

[15] F. Wirthmüller, M. Klimke, J. Schlechtriemen, J. Hipp, and M. Reichert, "Predicting the time until a vehicle changes the lane using LSTM-based recurrent neural networks," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 2021.

[16] Y. Hu, W. Zhan, and M. Tomizuka, "Probabilistic prediction of vehicle semantic intention and motion," in *IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 307–313.

[17] D. Xie, Z.-Z. Fang, B. Jia, and Z. He, "A data-driven lane-changing model based on deep learning," *Transportation Research Part C Emerging Technologies*, vol. 106, pp. 41–60, 2019.

[18] O. De Candido, M. Binder, and W. Utschick, "An interpretable lane change detector algorithm based on deep autoencoder anomaly detection," in *IEEE Intelligent Vehicles Symposium (IV)*, 2021, pp. 516–523.

[19] R. Chalapathy and S. Chawla, "Deep learning for anomaly detection: A survey," *CoRR*, vol. abs/1901.03407, 2019. [Online]. Available: http://arxiv.org/abs/1901.03407

[20] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, "Deep learning for anomaly detection: a review," *ACM Computing Surveys*, vol. 54, no. 2, 2021.

[21] O. Sharma, N. Sahoo, and N. Puhan, "Recent advances in motion and behavior planning techniques for software architecture of autonomous vehicles: A state-of-the-art survey," *Engineering Applications of Artificial Intelligence*, vol. 101, p. 104211, 2021.

[22] Q. Deng and D. Söffker, "Improved driving behaviors prediction based on fuzzy logic-hidden markov model (fl-hmm)," in *IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 2003–2008.

[23] L. Tong, S. Shi, and X. Gu, "Naive Bayes classifier based driving habit prediction scheme for VANET stable clustering," *Mobile Networks and Applications*, vol. 25, 2020.

[24] Y. Ma, S. Song, L. Zhang, L. Xiong, and J. Chen, "Lane change analysis and prediction using mean impact value method and logistic regression model," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021, pp. 1346–1352.

[25] F. Wirthmüller, J. Schlechtriemen, J. Hipp, and M. Reichert, "Teaching vehicles to anticipate: A systematic study on probabilistic behavior prediction using large data sets," *IEEE Transactions on Intelligent Transportation Systems (ITSC)*, 2020.

[26] D. Lee, Y. P. Kwon, S. McMains, and J. K. Hedrick, "Convolution neural network-based lane change intention prediction of surrounding vehicles for ACC," in *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2017, pp. 1–6.

[27] S. Ramyar, A. Homaifar, A. Karimoddini, and E. Tunstel, "Identification of anomalies in lane change behavior using one-class svm," in *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2016, pp. 004 405–004 410.

[28] P. Fan, Y. Liu, J. Guo, Y. Wang, and M. Qiu, "Anomalous state recognition of lane-changing behavior using a hybrid autoencoder architecture," in *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2020, pp. 1–6.

[29] H. Guo, K. Xie, and M. Keyvan-Ekbatani, "Lane change detection using naturalistic driving data," in *IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, 2021, pp. 1–6.

[30] J. Erdmann, "Lane-changing model in SUMO," in *Proceedings of the SUMO2014 Modeling Mobility with Open Data*, vol. 24, 2014.

[31] A. Benterki, M. Boukhnifer, V. Judalet, and M. Choubeila, "Prediction of surrounding vehicles lane change intention using machine learning," in *IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS)*, vol. 2, 2019, pp. 839–843.

[32] X. Liu, N. Masoud, and Q. Zhu, "Impact of sharing driving attitude information: A quantitative study on lane changing," in *IEEE Intelligent Vehicles Symposium (IV)*, 2020, pp. 1998–2005.

[33] C. Hydén, "Traffic conflicts technique: state-of-the-art," *Traffic safety work with video processing*, vol. 37, pp. 3–14, 1996.