ASME Journal of Autonomous Vehicles and Systems
Online journal at:
https://asmedigitalcollection.asme.org/autonomousvehicles

Feedforward Mutual-Information Anomaly Detection: Application to Autonomous Vehicles

Sasha M. McKee
Dynamics and Autonomous Robotics Laboratory,
Department of Mechanical Engineering and
Robotics Center,
University of Utah,
Salt Lake City, UT 84112
e-mail: sasha.mckee@utah.edu

Osama S. Haddadin
L3Harris Technologies,
Salt Lake City, UT 84116
e-mail: osama.s.haddadin@L3Harris.com

Kam K. Leang[1]
Fellow ASME
Dynamics and Autonomous Robotics Laboratory,
Department of Mechanical Engineering and
Robotics Center,
University of Utah,
Salt Lake City, UT 84112
e-mail: kam.k.leang@utah.edu

This paper describes a mutual-information (MI)-based approach that exploits a dynamics model to quantify and detect anomalies for applications such as autonomous vehicles. First, the MI is utilized to quantify the level of uncertainty associated with the driving behaviors of a vehicle. The MI approach handles novel anomalies without the need for data-intensive training; and the metric readily applies to multivariate datasets for improved robustness compared to, e.g., monitoring vehicle tracking error. Second, to further improve the response time of anomaly detection, current and past measurements are combined with a predictive component that utilizes the vehicle dynamics model. This approach compensates for the lag in the anomaly detection process compared to strictly using current and past measurements. Finally, three different MI-based strategies are described and compared experimentally: anomaly detection using MI with (1) current and past measurements (reaction), (2) current and future information (prediction), and (3) a combination of past and future information (reaction–prediction) with three different time windows. The experiments demonstrate quantification and detection of anomalies in three driving situations: (1) veering off the road, (2) driving on the wrong side of the road, and (3) swerving within a lane. Results show that by anticipating the movements of the vehicle, the quality and response time of the anomaly detection are more favorable for decision-making while not raising false alarms compared to just using current and past measurements.
[DOI: 10.1115/1.4064519]

Keywords: anomaly detection, mutual information, autonomous vehicles, intelligent decision making, transparency and trust

1 Introduction

The deployment of autonomous vehicles can help reduce driving-related fatalities and improve road safety [1]. However, safe integration of this emerging technology requires robust perception and decision-making in uncertain and complex environments. Currently, limitations in existing sensing and perception technologies lead to undesirable outcomes [2]. Recent work to address some of the challenges include incorporating a layer of safety through deep neural networks for prediction; unfortunately, large training datasets are required which can lead to poor robustness to new and uncertain situations [3]. Herein, a mutual-information (MI)-based anomaly quantification and detection process with reactive and predictive characteristics is described. This new scheme can be used to monitor when unusual behavior occurs so that dangerous behavior can be identified and actions can be taken to minimize or prevent accidents. The method is applied to autonomous driving in mobile robots to demonstrate efficacy, as visualized in Fig. 1.

Anomaly detection is a data-driven process that determines which event does not follow an expected trend [4]. The process of quantifying and detecting anomalies is needed in a wide range of applications, for instance, network monitoring [5], sensor-health monitoring [6], and autonomous driving [7]. Researchers have explored many different techniques for anomaly detection [8], including support vector machine [9]. In Ref. [10], a Bayesian network is utilized to find outliers in sensor data or to detect cyber and physical attacks in cyber-physical systems [11]. Clustering techniques have been employed, including k-means [12], nearest-neighbor [13], and multiview [14]. One drawback with clustering techniques is that there must be significant separation between normal and anomalous instances in the feature space, thus performance cannot be guaranteed [4]. Probabilistic-based approaches have been used to model normal behavior with Gaussian mixture models [15], non-parametric histograms [16], and kernel-density estimators [17]. However, since many probabilistic techniques assume a normal distribution [18], outside of this regime, the anomaly detection will perform poorly. Machine learning methods that involve neural networks (NN) [19,20] have been explored. In Ref. [21], the Kullback–Leibler divergence was incorporated as an additional input to a neural network. Even though many approaches have been introduced for anomaly detection, the majority of them, specifically, classification and artificial intelligence (AI), require large labeled datasets and are not robust to novel anomalies [22,23]. Large datasets can be unrealistic to obtain or may not represent the actual behavior well [24].

More recently, information-theoretic (IT) techniques have been exploited for anomaly detection. This framework models behavior

[1]Corresponding author.

Manuscript received November 7, 2023; final manuscript received January 18, 2024; published online February 13, 2024. Assoc. Editor: Vladimir Vantsevich.

Fig. 1 Anomaly detection based on vehicle behavior for self-driving vehicles

as a likelihood and when the likelihood is low, an anomaly is present. An advantage of IT techniques is that they can encode nonlinear behaviors [25] and work well even with imbalanced datasets, which is a challenge in many other techniques. Moreover, IT-based methods require little human involvement compared to other techniques [8]. However, one challenge is the IT framework can be sensitive to changes in the observed pattern [26]. Within the IT space, various forms of entropy have been proposed to quantify anomalies [27,28], such as energy measures [24] and Rényi entropy [26]. Revised and approximated forms of MI have been used [25,29] and combined with deep NN [23]. Many of the works that focus on MI only deal with selecting feature variables for large datasets [30,31], find most informative features [32], fuse multiple variables [33], act as an input to neural networks [6,34], and incorporate categorical data [35]. In contrast, this paper specifically uses MI to quantify and detect anomalies, which has not been thoroughly studied or considered in the past [36]. The closest work on using the information metric to quantify anomaly is the use of the Rényi information on network traffic, where the computation was performed offline [26]. To the best knowledge of the authors, this work is the first to exploit Shannon's mutual information to quantify and detect anomalies for real-time applications such as teasing out when undesirable behaviors occur in driving. Furthermore, implementation of existing IT approaches requires current and past measurements, and thus the detection process can exhibit delay [37] and be problematic for in-the-moment corrective measures. In this paper, a feedforward component is utilized to predict vehicle behavior to improve anomaly detection response time.

More specifically, this work uses MI to quantify the uncertainty between the actual and expected behaviors of the vehicle. This information-theoretic framework can be applied to multivariate datasets through the common representation of information. This approach makes no assumptions about the underlying process. Since a model of the expected behavior is provided, e.g., through a map, this method is robust to novel anomalies. By incorporating a predictive (feedforward) component through the dynamics model of the vehicle, the delay in calculating the MI is minimized compared to prior works using IT [26]. Many anomaly detection applications require quantifying anomalies in a timely manner [8] or predicting anomalies just moments before they occur so that evasive and defensive maneuvers in autonomous vehicles can be implemented effectively. The contributions of this paper are as follows:

(1) Using mutual information to quantify and detect anomalies;
(2) Improving the response time of anomaly detection through a predictive component involving the vehicle dynamics model;
(3) Studying the performance across three different MI-based approaches that involve reactive and predictive characteristics; and
(4) Validating the approach using physical experiments involving example autonomous mobile robots acting as vehicles.

2 Mutual-Information Anomaly Detection

The anomaly detection process comprises of two main steps: (1) predicting the behavior of the vehicle through the vehicle dynamics model and (2) quantifying anomalous behavior by mutual information, as shown in Fig. 2. An illustrative example of autonomous vehicles driving on a roadway is used as it demonstrates the basic

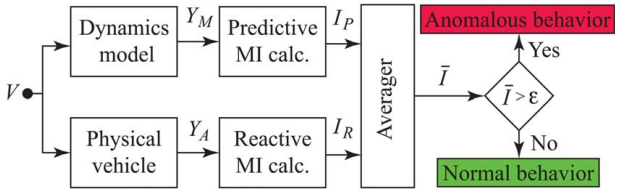

Fig. 2 Anomaly detection process that combines the vehicle's measured state (current and past), Y_A, and predicted state, Y_M. Reactive and predictive outputs are used to calculate the mutual information quantities. An averaging process creates the resultant mutual information $\bar{I}$, which is compared to some threshold, ε, for detecting anomalies.

capabilities of the algorithm in terms of characterizing when vehicles act anomalously. This method can be applied to other applications, such as network intrusion detection, by considering the relevant variables [36].

2.1 Measured and Predicted Vehicle Behaviors. Figure 3 shows the basic configuration of an autonomous vehicle moving in a 2D plane, where the state of the vehicle is defined by its pose (position and orientation). In particular, the vehicle's position is defined by the Cartesian coordinates of the center of mass of the vehicle, $l = (x, y) \in L$. The orientation (heading) is the vehicle's yaw angle $\phi = \psi \in \Phi$, measured relative to the world-frame X-axis. Using on-board sensors, it is assumed that current and past position and heading measurements are readily available. The measured heading direction (ϕ) is used to find an expected direction of travel, θ, given a map $\mathcal{M}$.

The vehicle's predicted behavior, Y_M, is created by a vehicle dynamics model given the input velocity V (see Fig. 2). The model can take various forms, but for simplicity and to illustrate the underlying concepts, the model is assumed to be linear with the form

$$\frac{Y_M(s)}{V(s)} = T(s) \tag{1}$$

where $V(s)$ denotes the input velocity, $Y_M(s)$ is the output of the model (e.g., vehicle location and orientation), and $T(s)$ is the transfer function that relates the input to output in the Laplace domain.

The anomaly detection process exploits the concept of mutual information to compare the trajectory of the vehicle to its expected trajectory based on the map $\mathcal{M}$. The MI will be formulated as a function of the measured past and current behaviors, and, herein, the predicted behavior from the dynamics model is further incorporated to improve the time response. For example, the reactive component of the mutual information at time-step k considers current and past time $t \in [k-n, \ldots, k-2, k-1, k]$ measurements, where $n \in \mathbb{N}$ denotes the number of past measurements. Likewise, the predictive component of the mutual information at time-step k

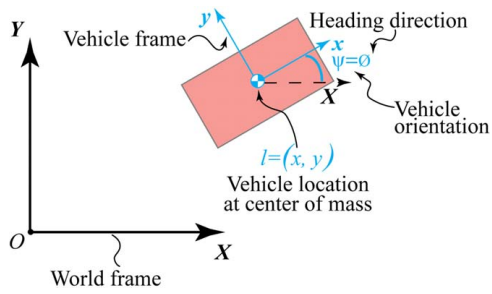

Fig. 3 Schematic of vehicle and coordinate frames. The vehicle states, [x, y, ψ], and measurements of location, l, and heading direction, ϕ.

considers future behaviors over the time window $t \in [k, k+1, k+2, \ldots, k+m]$, where window $m \in \mathbb{N}$ denotes the number of samples in future time. In addition, the complete time window of consideration is $w = [k-n, \cdot, k-2, k-1, k, k+1, k+2, \ldots, k+m]$. Later on, this time window is used to determine the average MI, denoted by $\bar{I}$, which is a function of the reactive MI, denoted by I_R, and predictive MI, denoted by I_P (see Fig. 2).

2.2 Quantifying Anomaly Using Mutual Information. Mutual information describes the amount of information associated with a random variable given the outcome of an event [38]. MI is used to quantify the vehicle's expected behavior (based on a map) with respect to the vehicle's heading direction. Leveraging Shannon's entropy (a measure of uncertainty), the mutual information between the vehicle's direction of travel and its location is

$$I(\Theta; L) = H(\Theta) - [H(\Theta, L) - H(L)] \tag{2}$$

where $H(\Theta)$ and $H(L)$ are the entropy of the expected direction of travel and vehicle location, respectively, and $H(\Theta, L)$ is the joint entropy.

A challenge with Eq. (2) is determining the entropy terms and corresponding probabilities. First, the entropy term $H(\Theta)$ is calculated by

$$H(\Theta) = -\sum_{i=k}^{k-n} p(\theta_i) \log_2 (p(\theta_i)) \tag{3}$$

where $p(\theta_i)$ represents the probability of the expected angle, θ_i, for the ith index, and k represents the current time instance. The entropy for the location, $H(L)$, is found in a similar manner, e.g.,

$$H(L) = -\sum_{i=k}^{k-n} p(l_i) \log_2 (p(l_i)) \tag{4}$$

where $p(l_i)$ is the probability of the location. Finally, the joint entropy, $H(\Theta, L)$, is given by

$$H(\Theta, L) = -\sum_{i=k}^{k-n} p(\theta_i, l_i) \log_2 (p(\theta_i, l_i)) \tag{5}$$

where $p(\theta, l) = p(\theta)p(l)$ is the joint probability for both the expected direction of travel and location (based on a map).

The equations above require that the probabilities be calculated. For example, the probability of the location is found using the current set of measurements, L, by

$$p(l) = \frac{a}{\Omega_L} \tag{6}$$

where $\Omega_L = unique(L)$ is the location vector L with non-repeating values, and a is the number of repetitions for the corresponding measurement [37]. The probability of the expected direction given the vehicle's heading is found as follows. First, the likely heading angles based upon the measured direction of travel, ϕ, is found by

$$p(\phi) = \mathcal{N}(E, \sigma_\phi, \mu_\phi) \tag{7}$$

where $\mathcal{N}$ represents a Gaussian distribution defined by the standard deviation, σ_ϕ, and mean, μ_ϕ. Additionally, the error is $E = \phi - \lambda$, where $\lambda \in \Lambda$ is the possible heading angle given the vehicle dynamics. Next, the expected direction(s) of travel, $\gamma \in \Gamma$, is found by the measured vehicle location, l,

$$p(\gamma) = \mathcal{F}(l, \mathcal{M}, \sigma_\mathcal{M}, \mu_\mathcal{M}) \tag{8}$$

which is a function of a given map, $\mathcal{M}$, and associated standard deviation, $\sigma_\mathcal{M}$ and mean, $\mu_\mathcal{M}$. Finally, the probability of the expected direction is determined by combining the likely heading angles and expected angles, e.g.,

$$p(\theta) = p(\phi)p(\gamma) \tag{9}$$

Intuitively, Eq. (9) describes the probability associated with an expected direction of travel given the measured heading angle, ϕ.

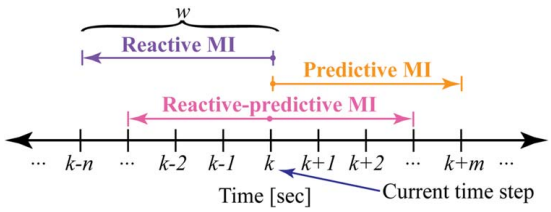

Fig. 4 Time window, w, for calculating the average MI, $\bar{I}$, for reactive, predictive, and reactive–predictive MI

2.3 Averaging Process. A collection of the MI values is combined to determine a moving average value, $\bar{I}$, which is a function of the predictive and reactive mutual information metrics, I_P and I_R (using Eq. (2)). The average MI, $\bar{I}$, is defined as

$$\bar{I}(I_R; I_P) = \eta\left(\frac{1}{n}\sum_{i=k}^{k-n} I_R\right) + \zeta\left(\frac{1}{m}\sum_{i=k}^{k+m} I_P\right) \tag{10}$$

where η and ζ are weighting terms and $\eta + \zeta = 1$. The average is determined over the desired time window, w, as illustrated in Fig. 4. The time window will be varied to explore the impact of the reactive versus predictive MI as well as sliding the time window relative to the current time-step k (see Fig. 4). An anomaly is detected when the average MI exceeds the threshold ε.

3 Experimental Setup

3.1 Overview. It has been shown that human error is the major cause in car accidents. In fact, the leading behaviors are driving: at high-speed, under the influence (DUI), while distracted, and aggressively [39]. To replicate these scenarios and to characterize the performance of the anomaly detection process, the following cases are considered:

- Case 1 (C1): Veering off-road, replicates distracted driving;
- Case 2 (C2): Wrong-way vehicle, represents aggressive driving by illegally passing a vehicle; and
- Case 3 (C3): Swerving within a lane, mimics DUI.

These cases will be implemented on a simple single roadway intersection as shown in Fig. 5. Vehicle behavior within the intersection conveniently demonstrates anomalous behavior.

For each case, five algorithms ($A1$–$A5$) will be compared as described in the following. The chosen time window is $w = 5$ s (see Fig. 4), which is approximately ten times longer than the settling time of the vehicle dynamics (see Sec. 3.3). The average MI that quantifies anomaly is calculated as follows:

($A1$) Reactive MI, where $\eta = 1.0$, $\zeta = 0.0$, $n = w$, and $m = 0$;
($A2$) Predictive MI, where $\eta = 0.0$, $\zeta = 1.0$, $n = 0$, and $m = w$;
($A3$) Reactive–predictive MI 1, where $\eta = 0.5$, $\zeta = 0.5$, $n = \frac{3}{5}w$, and $m = \frac{2}{5}w$;
($A4$) Reactive–predictive MI 2, where $\eta = 0.5$, $\zeta = 0.5$, $n = \frac{1}{2}w$, and $m = \frac{1}{2}w$; and
($A5$) Reactive–predictive MI 3, where $\eta = 0.5$, $\zeta = 0.5$, $n = \frac{2}{5}w$, and $m = \frac{3}{5}w$.

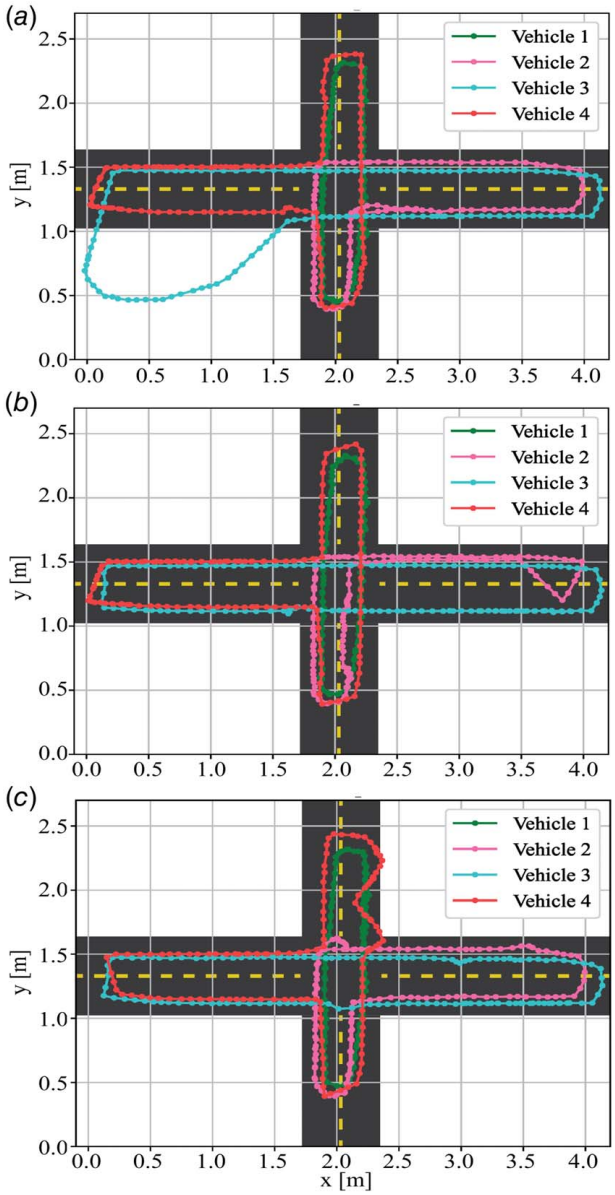

Fig. 5 Roadway intersection showing the measured trajectories of four different vehicles: (*a*) case 1 corresponds to veering off the road (vehicle 3), (*b*) case 2 shows a wrong-way driving vehicle (vehicle 2), and (*c*) case 3 involves a vehicle that swerves within its lane (vehicle 4)

In addition, at the start of the experiments, when k is outside of the time window, the average MI has $k-w$ instances of zeros.

3.2 Tracking Error for Anomaly Detection. A typical proxy that is often used to detect an anomaly is the vehicle tracking error, such as the Euclidean distance to the nearest road. This basic metric will be used as a basis for comparison (i.e., ground-truth tracking error). The Euclidean distance is defined as

$$d = \sqrt{(x_k - x_{\mathcal{M}})^2 + (y_k - y_{\mathcal{M}})^2} \tag{11}$$

where (x_k, y_k) are the Cartesian coordinates of the vehicle's center-of-mass at time-step k and $(x_{\mathcal{M}}, y_{\mathcal{M}})$ are the Cartesian coordinates to the center of the nearest road. The error value is then converted to an equivalent average MI value through a threshold, r, corresponding to a radius around the vehicle. In particular, the

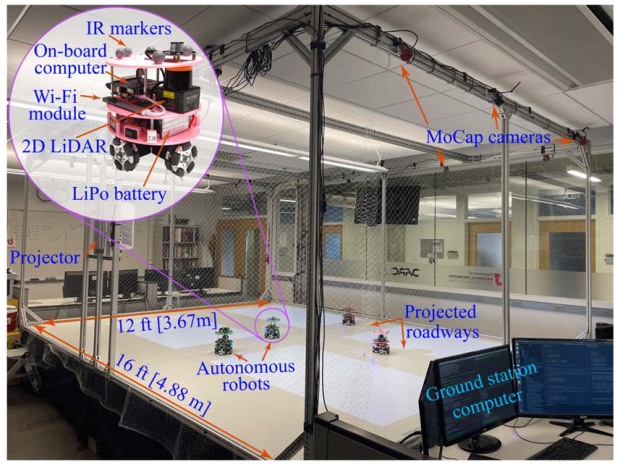

Fig. 6 The experimental platform equipped with motion capture (MoCap) infrared cameras and a ground station computer. The autonomous vehicles are custom-built and controlled on-board given a trajectory from the ground station. An overhead projector is used to visualize the road intersection during experiments.

pseudo-average MI, $\bar{I}_E$, associated with the tracking error is

$$\bar{I}_E = \begin{cases} 0 & d \le r \\ 3.2 & d > r \end{cases} \tag{12}$$

where the values 0 and 3.2 bits correspond to the minimum and maximum MI for the application of interest, respectively. This pseudo-average MI for the tracking error is compared to the output of the proposed MI-based algorithms.

3.3 Hardware Setup. The experimental test platform is shown in Fig. 6. The test platform houses a flat driving surface that is 16 feet by 12 feet (4.9 m × 3.7 m). The entire test volume is equipped with a total of 18 Optitrack Flex13 motion capture (MoCap) infrared cameras, operating at 120 Hz frame rate. A ground station computer, running Ubuntu 18.04, and the robot operating system (ROS) Melodic, is connected to the MoCap system and used to monitor the behavior of each vehicle, perform motion

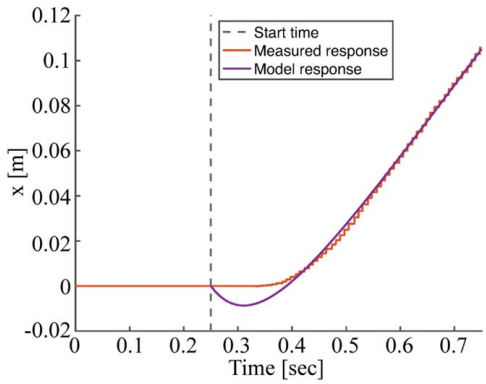

Fig. 7 Time response for the position (along the *x*-axis) of the vehicle, comparing the measured and modeled behavior

Table 1 Estimated gain and time delay of dynamics model

Gain (m/s)		Time delay (s)	
α_x	0.323	β_x	0.175
α_y	0.305	β_y	0.195
α_ψ	4.900	β_ψ	0.140

control, and calculate mutual information using measurements and predicted vehicle behavior to quantify anomalies.

As shown in the inset in Fig. 6, custom-built robotic vehicles were used in the experiments. Each vehicle carries an Odroid XU4 single-board computer (SBC), running Ubuntu 18.04 operating system, and ROS Melodic. The SBC is built on an A7 Octa-core CPU with 2 GB of LPDDR3 RAM. The SBC is interfaced with Robotis U2D2 to control the motors on the vehicle. A 4S 1.5-Ah lithium-polymer (Li-Po) battery is used to power the SBC and motors. Communication between each vehicle and the ground station is through a 2.4/5 GHz WiFi module. In addition, the vehicles are position-controlled with a discrete proportional-integral-derivative (PID) control loop and potential field is used for obstacle avoidance. The motion-control algorithms were developed under the ROS framework and ran on the SBC. A sequence of waypoints is sent to each vehicle from the ground station at approximately 100 Hz, and the on-board position controller continuously tracks the waypoints. Finally, all mutual information calculations

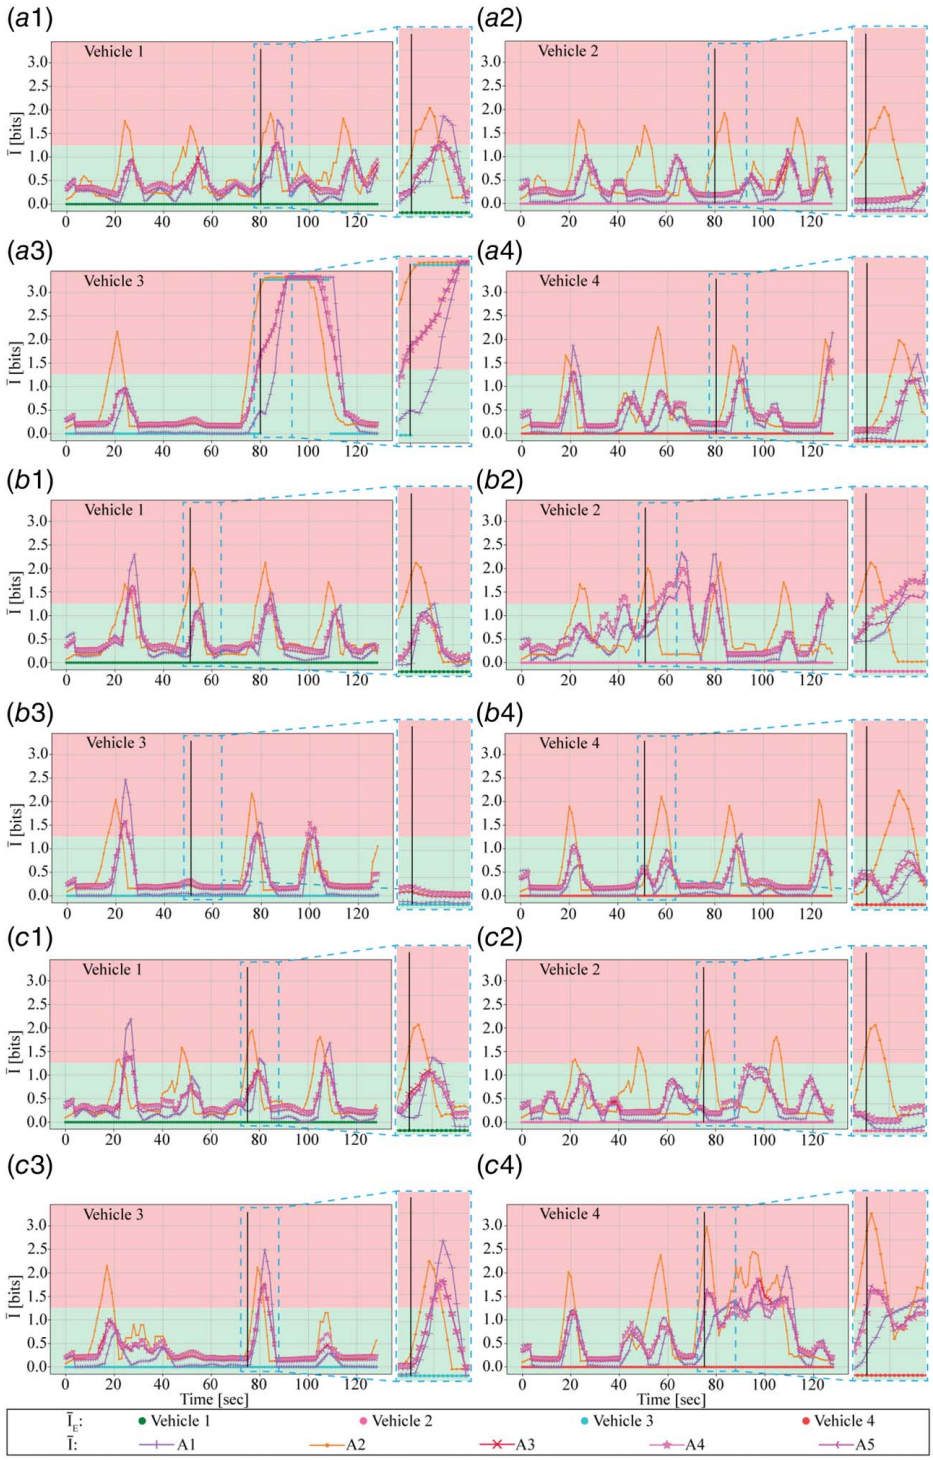

Fig. 8 Experimental results of the average mutual information, $\bar{I}$, versus time for (*a*) case 1—veering off-road with vehicle 3 acting as the anomaly, (*b*) case 2—wrong-way driving where vehicle 2 is the anomaly, and (*c*) case 3—swerving within lane with vehicle 4 acting anomalously. A solid vertical line indicates the anomalous activity start time.

Table 2 Results comparing the success of each algorithm averaged over all scenarios

Average algorithm success			
Algorithm	C (%)	FP (%)	FN (%)
Tracking error	95.93	0.65	3.42
A1 (Reactive)	92.57	5.36	2.07
A2 (Predictive)	84.30	14.02	1.68
A3 (Reac.–Pred.)	95.35	3.23	1.42
A4 (Reac.–Pred.)	95.23	3.17	1.62
A5 (Reac.–Pred.)	95.28	3.23	1.49

are conducted on the ground station; however, they can also be implemented in a distributed manner on board each vehicle.

A model of the vehicle dynamics is obtained by curve-fitting the measured step response, where the input is a desired vehicle velocity V. Figure 7 shows the measured vehicle time response in the x-direction as an illustrative example. The responses in the other two directions are similar and they are omitted for brevity. By the nature of the response shown in Fig. 7, an appropriate model of the vehicle dynamics consists of an integrator with a first-order Padé time delay, hence

$$\frac{Y_M(s)}{V(s)} = T(s) = \frac{\alpha}{s} e^{-\beta s} \approx \frac{\alpha}{s}\left(\frac{\frac{-\beta}{2}s + 1}{\frac{\beta}{2}s + 1}\right) \tag{13}$$

The transfer function model (Eq. 13) relates the input velocity $V(s)$ to the vehicle displacement $Y_M(s)$. The time-delay term is given by β [40]. The model parameters, α and β, for each vehicle, were found using a linear least squares fit of the measured response. Table 1 lists the estimated model parameters. As shown in Fig. 7, the proposed model captures the dominant dynamics of the vehicle with good accuracy. It is pointed out that more complex and sophisticated models of the vehicle dynamics, including accounting for wheel slip, steering dynamics, etc., can be used. In general, the use of a simplified model in this work serves to illustrate the basic application of the concepts proposed. In particular, the model is used to predict the behavior of the vehicle such that the MI process incorporates a component of anticipation (prediction) to improve transient response. Other approaches to predicting vehicle behavior can easily be incorporated into this MI-based anomaly detection.

The sample frequency used for the experiments is 10 Hz, which is sufficient given the dynamics of the vehicles of interest.

4 Results and Discussion

The experimental results are presented in Figs. 5 and 8, and Table 2. First, Fig. 5 shows the measured trajectories of four different vehicles (vehicle 1, vehicle 2, vehicle 3, and vehicle 4) as they traverse the roadway intersection. Each subplot represents a different case. Specifically, Fig. 5(*a*) shows results for case 1 where vehicle 3 veers off the road; Fig. 5(*b*) is for case 2 that demonstrates vehicle 2 driving on the wrong side of the road; and finally Fig. 5(*c*) is for case 3 where vehicle 4 swerves within its lane. All other vehicles are operating normally. Next, Fig. 8 shows the corresponding average mutual information over time for all the cases presented in Fig. 5. The outputs of the five different MI algorithms (i.e., *A*1, *A*2, *A*3, *A*4, and *A*5) for each case is plotted. Additionally, the MI version of the tracking error (ground truth), calculated by Eq. (12) with radius of 0.19 m, is also plotted. Furthermore, the threshold for detecting an anomalous event is $\varepsilon = 1.25$ bits. This is indicated on the plots by the two shaded regions where the lower section (green) is normal behavior and the upper (red) is anomalous. In addition, a solid black vertical line indicates the time at which the vehicle acts anomalously. More specifically, for case 1, vehicle 3 starts to act anomalously at $t = 80$ s. For case 2, vehicle 2 starts to act anomalously at $t = 51$ s. Finally, for case 3, vehicle 4 starts to act anomalously at $t = 75$ s. It should be noted that none of these anomalies are known to any of the algorithms or used for training of any kind, and thus, all results are novel anomalies. Table 2 lists the success rate for all algorithms averaged across all cases in terms of the percentage of correctness. In other words, quantifying the accuracy of detection by classifying each detection as follows: correct (C), a false positive (FP), which indicates an anomaly when the vehicle is not one, as well as a false negative (FN), which correlates to the detection of a normal vehicle when it is in fact an anomaly.

Examining the tracking error (ground truth) as a possible means for detecting anomalies, Fig. 8(*a*3) for case 1 shows that an anomaly was detected for vehicle 3. However, this same metric was unable to detect any anomalies for case 2 or 3, as shown in Figs. 8(*b*2) and 8(*c*4), respectively. This led to the tracking-error metric having the most amount of FN of 3.42% demonstrated in Table 2. By contrast, it was observed that the MI-based algorithms were able to detect all anomalies that were present. This result demonstrates the utility of the MI-based quantification compared to the tracking-error proxy. In fact, the MI-based approach presented can be used to encode additional behaviors to create a more robust composite metric for anomaly detection. Unfortunately, the tracking-error metric offers very limited capabilities in terms of anomaly detection.

In terms of time response for cases 1 and 3, most of the MI algorithms detected the anomaly at the instance that it starts occurring, e.g., see Figs. 8(*a*3) and 8(*c*4). The reactive algorithm (*A*1) lags by 6–9 s. In case 2, veering on the wrong side of the road, only the predictive technique, algorithm *A*2 identified the anomaly at the instance it occurred while algorithms *A*3 and *A*4 lagged behind by 5 s. The remaining algorithms *A*1 and *A*5 lagged behind by 7 s. Even if there is a lag, this technique is still capable of detecting all anomalies demonstrated, unlike the tracking-error metric. Furthermore, algorithms *A*3–*A*5, without any training, demonstrated an accuracy similar to the tracking-error metric of 95%, as shown in Table 2, while detecting various anomalies.

The results presented validated that mutual information can be used as an anomaly quantification and detection scheme. The approach is able to consider multiple events, e.g., vehicle location and heading. This technique did not require training and thus all anomalies detected in these scenarios are novel. Furthermore, by incorporating a vehicle model the predictive MI algorithms can generally detect the anomaly at the instance they start occurring, which is advantageous for real-time decision-making to minimize the impact of dangerous situations.

5 Conclusions

This paper described a novel mutual-information-based approach that incorporated a dynamics model to quickly and effectively quantify and detect anomalies for applications such as autonomous vehicles. The experimental results showed that across the five different MI calculations, the reactive–predictive approach demonstrated the best overall performance in terms of time response, correctness, and occurrences of false positive and negative. As expected, the reactive MI (A1) lagged behind by approximately 5 s, leading to the second-highest amount of false negative (2.07%) and false positive (5.36%). However, the purely predictive MI (*A*2) performed the worst overall with correctness of just 84.30%, and the highest false positive score (14.02%). Finally, the three reactive–predictive MI techniques (*A*3–*A*5) performed similarly with correctness near the tracking-error metric but were further able to detect more complex anomalies. Furthermore, this method requires no training and is robust to novel anomalies. However, a suitable model, that balances quality and computational demand, is needed for good accuracy. Finally, the mutual information framework with its ability to encode multiple variables of interest for improved robustness can be applied to other applications, such as network intrusion detection, which has a large imbalance of data.

Acknowledgment

This material is based upon work supported, in part, by the University of Utah and L3Harris Technologies. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

Data Availability Statement

The authors attest that all data for this study are included in the paper.

References

[1] Wang, H., Fan, R., Sun, Y., and Liu, M., 2022, "Dynamic Fusion Module Evolves Drivable Area and Road Anomaly Detection: A Benchmark and Algorithms," IEEE Trans. Cybernet., **52**(10), pp. 1–11.

[2] Castellano-Quero, M., Castillo-López, M., Fernández-Madrigal, J.-A., Arévalo-Espejo, V., Voos, H., and García-Cerezo, A., 2023, "A Multidimensional Bayesian Architecture for Real-Time Anomaly Detection and Recovery in Mobile Robot Sensory Systems," Eng. Appl. Artif. Intell., **125**(C), p. 106673.

[3] Stocco, A., and Tonella, P., 2022, "Confidence-Driven Weighted Retraining for Predicting Safety-Critical Failures in Autonomous Driving Systems," J. Softw. Evol. Process, **34**(10), p. e2386.

[4] Chandola, V., Banerjee, A., and Kumar, V., 2009, "Anomaly Detection: A Survey," ACM Comput. Surv., **41**(3), pp. 15:1–15:58.

[5] Callegari, C., Giordano, S., and Pagano, M., 2017, "Entropy-Based Network Anomaly Detection," IEEE International Conference on Computing, Networking and Communications, Silicon Valley, CA, Jan. 26–29, pp. 334–340.

[6] Lei, X., Xia, Y., Wang, A., Jian, X., Zhong, H., and Sun, L., 2023, "Mutual Information Based Anomaly Detection of Monitoring Data With Attention Mechanism and Residual Learning," Mech. Syst. Signal Process., **182**(4), p. 109607.

[7] Catal, O., Leroux, S., De Boom, C., Verbelen, T., and Dhoedt, B., 2020, "Anomaly Detection for Autonomous Guided Vehicles Using Bayesian Surprise," IEEE/RSJ International Conference on Intelligent Robots and Systems, Las Vegas, NV, Oct. 24–Jan. 24, pp. 8148–8153.

[8] Cook, A. A., Misirli, G., and Fan, Z., 2020, "Anomaly Detection for IoT Time-Series Data: A Survey," IEEE Int. Things J., **7**(7), pp. 6481–6494.

[9] Rajasegarar, S., Leckie, C., Palaniswami, M., and Bezdek, J. C., 2007, "Quarter Sphere Based Distributed Anomaly Detection in Wireless Sensor Networks," IEEE International Conference on Communications, Glasgow, UK, June 24–28, pp. 3864–3869.

[10] Safaei, M., Ismail, A. S., Chizari, H., Driss, M., Boulila, W., Asadi, S., and Safaei, M., 2020, "Standalone Noise and Anomaly Detection in Wireless Sensor Networks: A Novel Time-Series and Adaptive Bayesian-Network-Based Approach," J. Softw. Pract. Exper., **50**(4), pp. 428–446.

[11] Krishnamurthy, S., Sarkar, S., and Tewari, A., 2014, "Scalable Anomaly Detection and Isolation in Cyber-Physical Systems Using Bayesian Networks," ASME Dynamic Systems and Control Conference, San Antonio, TX, Oct. 22–24.

[12] Han, L., 2012, "Research of K-MEANS Algorithm Based on Information Entropy in Anomaly Detection," IEEE Conference on Multimedia Information Networking and Security, Nanjing, China, Nov. 2–4.

[13] Gu, X., Akoglu, L., and Rinaldo, A., 2019, "Statistical Analysis of Nearest Neighbor Methods for Anomaly Detection," Conference on Neural Information Processing Systems, Vancouver, Canada, Dec. 8–14.

[14] Ahmed, M., Mahmood, A. N., and Maher, M. J., 2015, "An Efficient Technique for Network Traffic Summarization Using Multiview Clustering and Statistical Sampling," EAI Endorsed Trans. Scalable Inf. Syst., **2**(5), pp. 1–9.

[15] Akouemo, H. N., and Povinelli, R. J., 2016, "Probabilistic Anomaly Detection in Natural Gas Time Series Data," Int. J. Forecast., **32**(3), pp. 948–956.

[16] Kind, A., Stoecklin, M. P., and Dimitropoulos, X., 2009, "Histogram-Based Traffic Anomaly Detection," IEEE Trans. Netw. Service Manage., **6**(2), pp. 110–121.

[17] Lang, C. I., Sun, F. -K., Lawler, B., Dillon, J., Dujaili, A. A., Ruth, J., Cardillo, P., Alfred, P., Bowers, A., McKiernan, A., and Boning, D. S., 2022, "One Class Process Anomaly Detection Using Kernel Density Estimation Methods," IEEE Trans. Semicond. Manuf., **35**(3), pp. 457–469.

[18] Chatterjee, A., and Ahmed, B. S., 2022, "IoT Anomaly Detection Methods and Applications: A Survey," Int. Things, **19**(3), p. 100568.

[19] Ali, M. H., Al Mohammed, B. A. D., Ismail, A., and Zolkipli, M. F., 2018, "A New Intrusion Detection System Based on Fast Learning Network and Particle Swarm Optimization," IEEE Access, **6**, pp. 20255–20261.

[20] Yao, H., Fu, D., Zhang, P., Li, M., and Liu, Y., 2019, "MSML: A Novel Multilevel Semi-Supervised Machine Learning Framework for Intrusion Detection System," IEEE Int. Things J., **6**(2), pp. 1949–1959.

[21] Wang, W., An, A., Zhang, Z., and Wang, Q., 2023, "Early-Warning of Generator Collusion in Chinese Electricity Market Based on Information Deep Autoencoding Gaussian Mixture Model," Electric Power Syst. Res., **221**(C), p. 109425.

[22] Ahmed, M., Naser Mahmood, A., and Hu, J., 2016, "A Survey of Network Anomaly Detection Techniques," J. Netw. Comput. Appl., **60**(C), pp. 19–31.

[23] Ahmad, Z., Khan, A. S., Nisar, K., Haider, I., Hassan, R., Haque, M. R., Tarmizi, S., and Rodrigues, J. J. P. C., 2021, "Anomaly Detection Using Deep Neural Network for IoT Architecture," Appl. Sci., **11**(15), p. 7050.

[24] Leach, M. J. V., Sparks, E. P., and Robertson, N. M., 2014, "Contextual Anomaly Detection in Crowded Surveillance Scenes," Pattern Recogn. Lett., **44**(C), pp. 71–79.

[25] Wang, Q., Shen, Y., and Zhang, J. Q., 2005, "A Nonlinear Correlation Measure for Multivariable Data Set," Phys. D: Nonlinear Phenom., **200**(3-4), pp. 287–295.

[26] Kopylova, Y., Buell, D. A., Huang, C.-T., and Janies, J., 2008, "Mutual Information Applied to Anomaly Detection," J. Commun. Netw., **10**(1), pp. 89–97.

[27] Gautam, S. K., and Om, H., 2015, "Anomaly Detection System Using Entropy Based Technique," IEEE 1st International Conference on Next Generation Computing Technologies, Dehradun, India, Sept. 4–5, pp. 738–743.

[28] Arackaparambil, C., Bratus, S., Brody, J., and Shubina, A., 2010, "Distributed Monitoring of Conditional Entropy for Anomaly Detection in Streams," IEEE International Symposium on Parallel & Distributed Processing, Workshops and Ph.D. Forum, Atlanta, GA, Apr. 19–23, pp. 1–8.

[29] Kay, S., and Emge, D., 2023, "Anomaly Detection Via Estimated Mutual Information and Its Relationship to the GLRT," IEEE Signal Process. Lett., **30**, pp. 220–223.

[30] Lima, C. F. L., de Assis, F. M., and Protásio, C. P. C., 2010, "Decision Tree Based on Shannon, Rényi and Tsallis Entropies for Intrusion Tolerant Systems," IEEE Fifth International Conference on Internet Monitoring and Protection, Barcelona, Spain, May 9–15, pp. 117–122.

[31] Hoque, N., Bhattacharyya, D. K., and Kalita, J. K., 2014, "MIFS-ND: A Mutual Information-Based Feature Selection Method," Exp. Syst. Appl., **41**(14), pp. 6371–6385.

[32] Kappaganthu, K., and Nataraj, C., 2011, "Feature Selection for Fault Detection in Rolling Element Bearings Using Mutual Information," ASME J. Vib. Acoust., **133**(6), p. 061001.

[33] Feng, C., Liu, C., and Jiang, D., 2023, "Unsupervised Anomaly Detection Using Graph Neural Networks Integrated With Physical-Statistical Feature Fusion and Local-Global Learning," Renew. Energy, **206**(C), pp. 309–323.

[34] Sheng, S., and Wang, X., 2023, "Network Traffic Anomaly Detection Method Based on Chaotic Neural Network," Alexandria Eng. J., **77**, pp. 567–579.

[35] Liu, D., Lung, C.-H., Seddigh, N., and Nandy, B., 2014, "Entropy-Based Robust PCA for Communication Network Anomaly Detection," IEEE/CIC International Conference on Communications in China, Shanghai, China, Oct. 13–15, pp. 171–175.

[36] Wenke, L., and Dong, X., 2001, "Information-Theoretic Measures for Anomaly Detection," IEEE Symposium on Security and Privacy, Oakland, CA, May 14–16, pp. 130–143.

[37] Waskita, A., Suhartanto, H., and Handoko, L. T., 2016, "A Performance Study of Anomaly Detection Using Entropy Method," IEEE International Conference on Computer, Control, Informatics and its Applications, Tangerang, Indonesia, Oct. 3–5, pp. 137–140.

[38] Shannon, C. E., 1948, "A Mathematical Theory of Communication," Bell Syst. Techn. J., **27**(3), pp. 379–423.

[39] National Highway Traffic Safety Administration, 2023, https://www.nhtsa.gov/risky-driving.

[40] Natori, K., 2012, "A Design Method of Time-Delay Systems With Communication Disturbance Observer by Using Pade Approximation," IEEE International Workshop on Advanced Motion Control, Sarajevo, Bosnia and Herzegovina, Mar. 25–27, pp. 1–6.