CPS Approach to Checking Norm Operation of a Brake-by-Wire System

Kyong-Tak Cho, Kang G. Shin
Real-Time Computing Laboratory
Dept. of EECS
University of Michigan – Ann Arbor, U.S.A
{ktcho, kgshin}@umich.edu

Taejoon Park
Real-Time CPS Research Laboratory
Dept. of Information and Communication Engineering
DGIST, Daegu, Republic of Korea
tjpark@dgist.ac.kr

ABSTRACT

For better controllability and energy-efficiency, more vehicle functions are being implemented via electronic control systems in place of traditional mechanical control systems. However, such transitions are creating new, unprecedented risks such as software bugs or hardware glitches, all of which can lead to serious safety risks. Recent real-world examples and research literature have been covering them under the name of *vehicle misbehavior*. In this paper, we present a new way of checking norm operations, called `BAD` (Brake Anomaly Detection), which detects any vehicle misbehavior in the Brake-by-Wire system. We focus on the braking system since it is a prototypical safety-critical and cyber-physical system. We first propose a new method for constructing norm models of braking and then show how anomalies are detected by `BAD` using the constructed models. Finally, we discuss how to verify the results, especially in the context of false positives. Our evaluation results show that `BAD` can effectively detect various types of anomaly in the braking system.

Categories and Subject Descriptors

C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems; C.4 [**Performance of Systems**]: Reliability, availability, and serviceability

General Terms

Algorithms, Design, Reliability, Verification

Keywords

Anomaly detection, vehicle misbehavior, braking system

1. INTRODUCTION

Advances in vehicle technology have made vehicles efficiently maneuverable. Sensors and drive-by-wire technologies – which replace the traditional mechanical control systems with electronic ones – improve response time, safety, control precision, and also reduce weight. Moreover, use of wireless communication technologies for Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) further improves safety, enhances traffic manageability, and offers many other benefits.

Do these bring only benefits to vehicles and their drivers/ passengers? Recent real-world examples and research literature answer "No" to this question, showing that new types of defects may lead to critical accidents.

As of January 2010, Toyota recalled approximately 7.5 million vehicles to resolve an "unintended acceleration" problem. At first, the reasons for this problem were thought to be a floor mat problem, accelerator pedal problem, or driver error of pedal misapplication. However, in October 2013, a jury ruled against Toyota and found that software deficiencies in the Throttle-by-Wire system could have caused the unintended acceleration. Barr [1] testified that a single bit flip or stack overflow can trigger unintended acceleration. Moreover, he stated that more than 16 million combinations of task deaths can occur in thousands of different states of the system, possibly causing an *unpredictable range* of vehicle misbehavior.[1] In other words, software bugs and deficiencies are now major problems in vehicles as they can cause vehicle misbehavior. Moreover, not only software but also hardware deficiencies can also cause the misbehavior [2].

The above examples reveal that, ironically, new in-vehicle electronic components deployed for enhanced controllability and safety may in fact expose drivers to more risk. The main risk factor being not human error but vehicle misbehavior stemming from software bugs or hardware glitches. Moreover, the unpredictable range of vehicle misbehavior would mean drivers or even manufacturers would not know *when* and *how* vehicles can go wrong. Thus, it is essential to analyze and detect such vehicle misbehavior in real time.

To address the unpredictability of vehicle misbehavior, we propose an anomaly detection method, called `BAD` (Brake Anomaly Detection) for the braking system. Since anomaly detection is to identify any observation which does not conform to an expected pattern or behavior, it is a natural candidate for detecting unpredictable misbehavior. We will focus on the Brake-by-Wire system, since it is the most important vehicle subsystem for safety.[2]

ICCPS '15 April 14 - 16, 2015, Seattle, WA, USA

http://dx.doi.org/10.1145/2735960.2735977

[1] Vehicle misbehavior refers to the vehicles being maneuvered differently from the driver's intention.

[2] Due to the safety-critical nature of braking, a fully electronic Brake-by-Wire system, which simply uses an electric motor and no hydraulic brake fluid, is yet to be deployed in passenger cars.

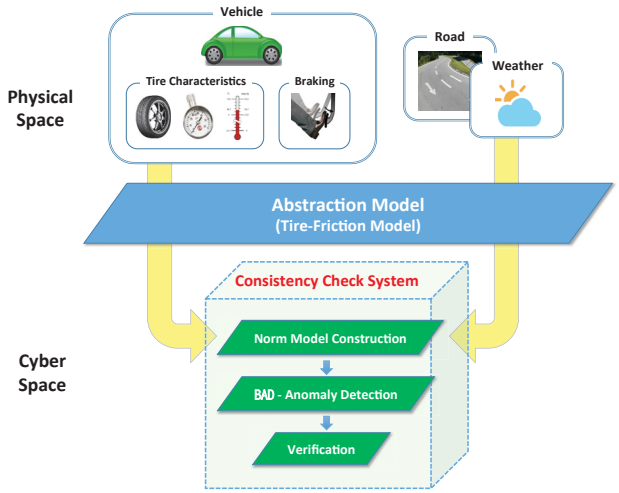

Figure 1: Cyber-physical approach to checking the norm operation of braking.

We take a *cyber-physical* approach by integrating and coordinating cyber and physical parts to detect braking anomalies as in Fig. 1. For the physical parts, there should be an understanding of the vehicle's braking characteristics and the driver's intention on braking. Moreover, other physical parts such as the road pavement condition and weather may also affect the braking characteristics. By taking them into account in the cyber space, we abstract the physical properties using a tire-friction model and design `BAD` using the abstracted model in detecting any braking anomalies.

The rest of the paper is organized as follows. Section 2 provides background of braking-related theories, and Section 3 formally states our problem. Section 4 presents the proposed consistency check system, and Section 5 describes how norm models can be constructed. Section 6 presents the details of `BAD`, and Section 7 describes the verification process. Section 8 evaluates the proposed system using realistic simulations, and finally, Section 9 concludes the paper.

2. BACKGROUND

While driving a vehicle, tire-road forces are the only forces that a vehicle experiences from the road. When a force is applied to the wheels during acceleration or braking, the tires instantaneously slip on the surface by a certain level and thus make the vehicle accelerate or decelerate. Tire-friction models describe the relationship between such level of slip and the applied tire traction forces.

2.1 Tire Friction Model

The tire-friction model describes the relationship between the tire slip and the traction force (i.e., slip-force relationship). The model can be used for traction control; for example, the tire slip level is used as an input for ABS activation.

Various tire-friction models describe the slip-force relationship in a numerical form. One simple model is a linear tire model which is commonly used for the vehicle bicycle model. On contrary to linear tire models, nonlinear tire models such as the Pacejka's tire model [3] and the Brush model [4] capture the effect of friction variations. These nonlinear tire-friction models have been shown to suitably match the field measured data [3–5, 7].

2.2 Slip Ratio

The amount of tire slip is represented by a metric called the *slip ratio*:

$$\sigma := \frac{r_{eff} w_w - V}{\max\left(r_{eff} w_w, V\right)}, \tag{1}$$

where r_{eff} is the rolling tire radius, w_w the wheel angular velocity, and V the vehicle speed. During acceleration, the instantaneous wheel speed, $r_{eff} w_w$ is higher than V, so $\sigma > 0$. On the other hand, the wheel speed drops below V during braking, leading to $\sigma < 0$. Due to the dynamics in driving, and noise and error in measurements, it is unlikely that the measured slip ratio is 0, i.e., $\sigma \neq 0$.

Since the slip ratio is derived from the speeds of the vehicle and wheels, this metric can be used to characterize the vehicle behavior. Details of how the slip ratio can be acquired through measurements are provided in Appendix A.

2.3 Normalized Traction Force

Normalized traction force is derived by dividing the applied longitudinal and lateral force on the tire by the normal vertical load on the tire. The longitudinal force is generated during braking or acceleration, while the lateral force is generated when cornering. Since our focus is on braking, we will just consider the longitudinal forces on tires. The normalized *longitudinal* force can thus be expressed as

$$\rho_x := F_x / F_z. \tag{2}$$

where F_x and F_z represent the longitudinal and normal forces acting on a tire, respectively. Since we only need to extract the longitudinal force of the tires, in the remainder of this paper, we will use the terms "normalized traction force" and "normalized longitudinal force", interchangeably.

The normalized traction force is proportional to the brake torque, which is, in turn, proportional to the brake pedal position. Thus, this metric captures the *driver's intention* on braking. Appendix B details the derivation of normalized traction force using in-vehicle sensor measurements.

2.4 Friction Coefficient

The maximum achievable traction force varies with the road pavement condition that a vehicle is running on. When running on roads with low friction (e.g., snowy roads), the maximum degree of achievable normalized traction force is lower than that for roads with high friction (e.g., dry roads). Such maximum degree is defined as the friction coefficient, μ, i.e., $\mu = \max|\rho_x|$. This implies that during acceleration, ρ_x cannot possibly exceed the road's friction coefficient, μ, whereas during braking, it cannot go below $-\mu$. Typically, μ ranges between 0 and 1, which varies in the range of 0.8∼1, 0.4∼0.7, and 0.1∼0.3 for dry, wet, and snowy/icy pavements, respectively.

The friction coefficient not only implies the limits of normalized traction force but also affects the slip-force relationship. More details on such a relationship will be provided when we discuss the Brush tire-friction model in Section 5.1. Since the friction coefficient is dependent on the road condition and irrelevant to the vehicle's characteristics, this metric can be considered as a representation of the *environment*.

3. PROBLEM STATEMENT

Due to the introduction of new technologies on vehicle subsystems, many critical issues such as software bugs or

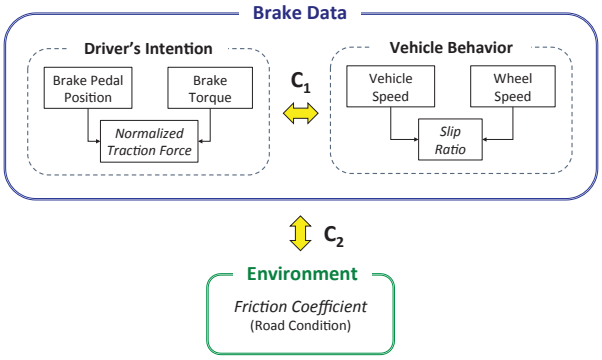

Figure 2: Checking consistency between the driver's intention and vehicle braking behavior, and brake data and environment.

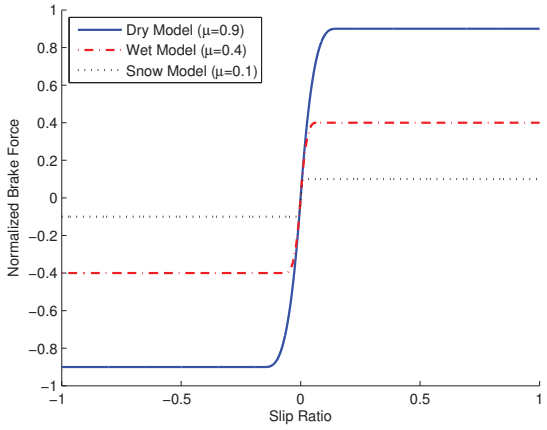

Figure 3: Norm braking models, which are based on the Brush tire-friction model and constructed by extrapolating low slip measurements.

hardware glitches are emerging and causing potential threats. They can occur for numerous reasons and scenarios, and can result in unpredictable braking misbehavior.

In the unintended acceleration example, there is a mismatch between the driver's intention to slow down the vehicle by pressing the brake pedal and the actual vehicle behavior which was acceleration.

Due to such new possibilities of unintended behavior, a vehicle should now be able to discover, via *consistency checking*, any vehicle behavior that is inconsistent with the driver's intention or the environment. Our goal in this paper is thus to construct a *real-time consistency check system* that can reliably detect, in real time, any abnormal brake data from the vehicle's braking system. Abnormal brake data refers to in-vehicle data that can trigger unintended and unpredictable vehicle misbehavior. We define brake data as a tuple of slip ratio and normalized traction force, (σ, ρ_x), each representing the vehicle behavior and the driver's intention, respectively.

Using the measured brake data during acceleration or braking, the norm model of braking can be constructed. By using the thus-constructed norm model, we want to detect any abnormal brake data that can lead to vehicle misbehavior. Further verifications on false-positive results are done using other in-vehicle sensor data.

4. CONSISTENCY CHECK SYSTEM

To detect vehicle misbehavior in braking, we propose a consistency check system as shown in Fig. 2. The vehicle's behavior should be consistent with the driver's intention as well as with the driving environment (e.g., friction coefficient). To meet these requirements, the system cross-checks the measured brake data in two different ways.

C1. *Driver's Intention vs. Vehicle Braking Behavior:* The system first checks whether the vehicle behavior is consistent with the driver's intention of braking. Through measurements on the brake pedal position and brake torque, the normalized traction force that represents the driver's intention on braking is determined. Moreover, the slip ratio value, which represents the vehicle behavior, is also acquired via in-vehicle sensors. By using a tire-friction model, which would be a reference for the norm model for braking, consistency check between the driver's intention (normalized traction force) and the vehicle's behavior (slip ratio) is performed.

C2. *Brake Data vs. Environment:* The system also checks whether the brake data, (σ, ρ_x), is in accordance with environmental conditions around the vehicle. As stated in Section 2.4, the achievable degree of traction force is dependent on, and limited by, the road's friction coefficient (typifying the driving environment). Exploiting such a fact, if the measured brake data yields a friction coefficient that is inconsistent with the current road condition or with the reality, we consider it as an indicator of vehicle misbehavior.

5. NORM MODEL CONSTRUCTION

To address difficulties in predicting the range of vehicle misbehavior, we take an anomaly detection approach, identifying any observations that do not conform to an expected pattern or behavior. Thus, for anomaly detection, a norm model describing the expected pattern or behavior should be constructed beforehand. Since tire-friction models represent the *norm* relationship between the slip ratio and the normalized traction force (i.e., a norm slip-force relationship), we use it as a basis for norm model construction.

5.1 Brush Tire-Friction Model

Among existing non-linear tire models, we choose the Brush model [4] to represent the norm slip-force relationship. It is a physics-based model and is shown to fit well with field measurement data. The model describes the slip-force relationship by constructing the model using two parameters: tire stiffness, C, and friction coefficient, μ. The tire stiffness represents the braking efficiency and is dependent on inflation pressure, tire wear, temperature, and tire type [5, 7]. In other words, tire stiffness is a parameter that reflects the vehicle characteristic in general and the tire characteristic in particular. In addition, the friction coefficient represents the limit of achievable traction and the road condition.

Given C and μ, the Brush model expresses the relationship

between the traction force, ρ_x and the slip ratio, σ as

$$\rho_x^{C,\mu}(\sigma) = \begin{cases} C\sigma + \frac{1}{3}\frac{C^2\sigma|\sigma|}{\mu} - \frac{1}{27}\frac{(C\sigma)^3}{\mu^2} & \text{if } |\sigma| < \frac{3\mu}{C} \\ \mu \cdot sign(\sigma) & \text{otherwise.} \end{cases} \quad (3)$$

Three different Brush models with the same C but different μ values are shown in Fig. 3 as an example. We define each model as the *norm braking model* for a specific road condition since it represents the norm slip-force relationship under a certain μ value. Moreover, we define the set of constructed norm braking models as a *norm model set*.

In the region where slip ratio is low (i.e., $|\sigma| \approx 0$), $\rho_x \approx C\sigma$, meaning that the slip-force relationship is linear with its slope equal to tire stiffness,[3] but is independent of the friction coefficient. This in turn means that, for a given vehicle, norm braking models with different μ values will all have the same slip-force relationship in the low slip region. Field measurements show that this relationship holds for $|\sigma| \leq \sigma_{linear}$, where typically $\sigma_{linear} \simeq 0.02$ [5,7,8]. We will refer to the region where $|\sigma| \leq \sigma_{linear}$ as the *low slip region*, and $|\sigma| > \sigma_{linear}$ as the *mid–high slip region*.

As the absolute slip ratio increases, the tire stiffness has less influence on the slip-force relationship. The friction coefficient then takes over that role and affects the shape of the slip-force curve. Therefore, the slip-force relationship becomes non-linear in the mid–high slip region. Accordingly, as shown in Fig. 3, the norm braking models have different values depending on the friction coefficient.

During ordinary driving, the absolute slip ratio rarely exceeds 1∼2% ($\simeq \sigma_{linear}$), and thus the slip-force relationship is usually linear [10]. Moreover, the tire stiffness is the same for both acceleration ($\sigma > 0$) and braking ($\sigma < 0$), i.e., the slip-force curve has a *symmetric* characteristic.

The benefits of using the Brush model to construct the norm braking model are as follows. First, since it only has two parameters to construct the model, it is easy to learn the parameters and construct the norm model. Second, even in normal driving scenarios when the absolute slip ratio is low (0∼2%), the entire slip-force curve can be constructed just by measuring brake data in the low slip region and extrapolating it outside the measurement region. Lastly, due to the symmetric characteristic of the slip-force curve between acceleration and braking, the norm model of braking can be constructed based on in-vehicle data while accelerating, and be used while braking.

5.2 Tire Stiffness Identification

Norm braking models can be constructed once the tire stiffness parameter of the Brush model is identified. The other parameter, μ, is then used as an input to construct different types of norm models for different pavement conditions as in Fig. 3. For example, a model with $\mu = 0.9$ can be interpreted as a dry road norm model, whereas a model with $\mu = 0.1$ can be interpreted as a snowy road norm model.

Exploiting the symmetrical characteristic of the norm slip-force curve, we can identify the tire stiffness, and accordingly construct the norm braking model during *acceleration*. This approach has the following advantages.

- Since the construction is based on acceleration data, any abnormal brake data generated during braking does no harm to the norm model construction.
- Most of the driving time is composed of acceleration rather than braking, and hence, there should be sufficient data to construct the norm model.
- Since acceleration and braking are performed in turn, not simultaneously, the norm model constructed during acceleration can be immediately applied to braking. This provides freshness in using the norm model.

[3]For these reasons, the tire stiffness is also called the *slip slope* [10].

Algorithm 1 Norm model construction and `BAD` in the low slip region

Initialize:
  $\varphi[0] = 0$
  $P[0] = \delta I$
function TSUPDATE(x, y, e) ▷ Update tire stiffness
  $K[n] \leftarrow \frac{\lambda^{-1}P[n-1]x[n]}{1+\lambda^{-1}x^T[n]P[n-1]x[n]}$
  $P[n] \leftarrow \lambda^{-1}(P[n-1] - K[n]x^T[n]P[n-1])$.
  $\varphi[n] \leftarrow \varphi[n-1] + K[n]e[n]$
  return $\varphi[n]$
end function
for n^{th} step **do**
  measure and derive $\rho_x[n]$ and $\sigma[n]$
  $x[n] \leftarrow \sigma[n]$
  $y[n] \leftarrow \rho_x[n]$
  if $|x[n]| \leq \sigma_{linear}$ **then** ▷ Linear region
    $e[n] \leftarrow y[n] - x[n]^T\varphi[n-1]$
    if x[n]>0 **then** ▷ Tire stiffness identification
      $C_{norm} \leftarrow$ TSUPDATE($x[n], y[n], e[n]$)
    else ▷ `BAD` via error monitoring
      if $|e[n]| > \Gamma_{anomaly}$ **then**
        go to Verification stage
      else
        $C_{norm} \leftarrow$TSUPDATE($x[n], y[n], e[n]$)
      end if
    end if
  end if
  Construct norm model set based on C_{norm}
end for

Since each vehicle would have different tire stiffness, the identification process should be individually performed to construct its own norm braking model. Considering the fact that the tire stiffness, C is the slip slope in the low slip region, we formulate the identification of tire stiffness as a linear parameter identification problem:

$$y[n] = x^T[n]\varphi[n] + e[n] \quad (4)$$

where $y[n] = \rho_x$ is the measured normalized traction force, $x[n] = \sigma$ the measured slip ratio, and $\varphi[n] = C$ the unknown tire stiffness parameter. Note that this formulation is only valid in the low slip region.

For identifying the unknown parameter C, we use the Recursive Least Square (RLS) algorithm [11], which considers the residual as a metric and aims to minimize the sum of the squares in the modeling errors.

Other identification algorithms such as Total Least Square (TLS) and Damped Least Square (DLS) can be used. In contrast to RLS, TLS considers the true Euclidean distance and performs an orthogonal linear regression, whereas DLS identifies the parameter based on non-linear regression. Due to such aspects of TLS and DLS, they might identify the tire stiffness with higher accuracy than RLS.

However, for the purpose of real-time anomaly detection, gain of such a higher identification accuracy from using TLS or DLS is offset by their high complexity. TLS requires an execution of the Singular Value Decomposition (SVD),

a non-recursive matrix decomposition, and thus would be computationally expensive to be used for real-time identification. Moreover, in some cases, DLS can only find the local minimum for curve fitting. Thus, to find the global minimum, a large number of iterations and careful determination of the search start point are required.

RLS algorithms are known to have a computation complexity of $\mathcal{O}(N^2)$ per time iteration, where N is the size of the data matrix. However, for our purpose of norm model construction, we only use a single brake data at each step, and thus incurs low computational complexity. We will later show through evaluation that RLS is sufficient for identifying the tire stiffness, and thus constructing norm models.

The tire stiffness identification process is shown in lines 9 to 15 in Algorithm 1. After initializing the unknown parameter $\varphi[n]$=0, and the covariance matrix $P[n]$ as $P[0] = \delta I$, where δ is a RLS parameter that determines the initial estimate of $P[n]$, the measured slip ratio, $x[n]$, and the normalized traction force, $y[n]$, is used for identification. First, the identification error is derived using the measurements at step n and the estimated parameter at step $n-1$ as in line 13. Then, by using a TSUPDATE function, the gain vector, $K[n]$, and the covariance matrix, $P[n]$, are updated so that the tire stiffness, $\varphi[n]$, can be identified. This procedure of tire identification continues iteratively during acceleration.

In the TSUPDATE function, we use an RLS parameter, λ, which is a forgetting factor. Since old information is less representative for the current status, by giving exponentially less weight with a factor of λ^{-1} to older error samples, their influence is reduced, providing a fresher result. Moreover, as samples get older, their effect on the result becomes negligible and can thus be discarded. As a result, only a limited amount of recent data is required to be stored in the system. A typical value of λ ranges from 0.9 to 1.

Once the tire stiffness of the vehicle, C_{norm}, is learned, by plugging in that and different friction coefficient values into Eq. (3), a norm model set consisting of various norm braking models can be constructed as shown in Fig. 3.

6. BAD – BRAKE ANOMALY DETECTION

Once the norm braking models are constructed, BAD can be performed. BAD is divided into two parts: one for the low slip region and the other for the mid–high slip region. The reason for this division is that the two regions show different characteristics of the slip-force curve's linearity. An overview of the proposed BAD algorithm is shown in Fig. 4, and elaborated in the following subsections.

6.1 Low Slip Region

In the low slip region, the slip-force relationship is not only linear but also independent of the friction coefficient. In other words, the brake data is independent of the environment, meaning no need to perform the consistency check, C2. BAD, therefore, only performs consistency check, C1, between the driver's intention (normalized traction force) and vehicle behavior (slip ratio).

By using the linear relationship between the two elements of brake data, BAD exploits the RLS algorithm in detecting abnormal brake data. We also numerically show the strength of BAD in the low slip region.

6.1.1 BAD *via Identification Error Monitoring*

The objective of BAD is to detect, if any, abnormal brake

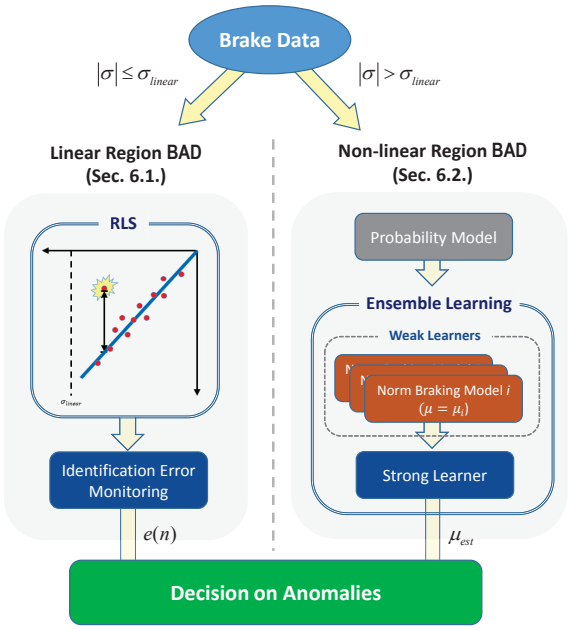

Figure 4: Overview of the proposed BAD (Brake Anomaly Detection).

data that deviates greatly from the norm braking models. Since the friction coefficient has negligible impact on the slip-force relationship in the low slip region, a single norm braking model is sufficient to represent all cases of norm behavior.

BAD in low slip region is shown in lines 16 to 22 in Algorithm 1, which basically runs the same RLS procedure with an additional step of monitoring the degree of identification error, $|e[n]|$, during braking. $e[n]$ can be considered as the residual which is not explained by the norm model. Thus, the magnitude of the residual, $|e[n]|$ is used as the anomaly score of the examined brake data. If the score does not exceed a pre-determined threshold, $\Gamma_{anomaly}$, the data is considered normal, and the tire stiffness parameter, C_{norm} is updated as in line 20. On the other hand, if it exceeds the threshold and shows a high anomaly score, the brake data might be an anomaly. To deal with possible false-positives, a verification process is performed.

6.1.2 *Properties of* BAD *in low slip region*

BAD achieves high detection capability and accuracy by *amplifying* the anomaly score caused by consecutive vehicle misbehavior. Unusually large changes in the brake data are also detected.

Let us consider a case where abnormal vehicle behavior has occurred at step n, where $y[n]$ highly deviates from the expected value of $\varphi[n-1]x[n]$. According to the update equations of $e[n]$ and $\varphi[n]$, the anomaly score at the next step, $n+1$, can be derived as

$$e[n+1] = \Delta y - \varphi[n-1]\Delta x + (1+\kappa)\, e[n], \qquad (5)$$

where $\Delta x = x[n+1] - x[n]$, $\Delta y = y[n+1] - y[n]$, and $\kappa = K[n]x[n+1]$. According to the update equation for $K[n]$, $\kappa > 0$. Due to the abnormal behavior at step n, $e[n]$ would be larger than usual. Referring to (5), this value gets amplified by $1+\kappa$ at the next step. In other words,

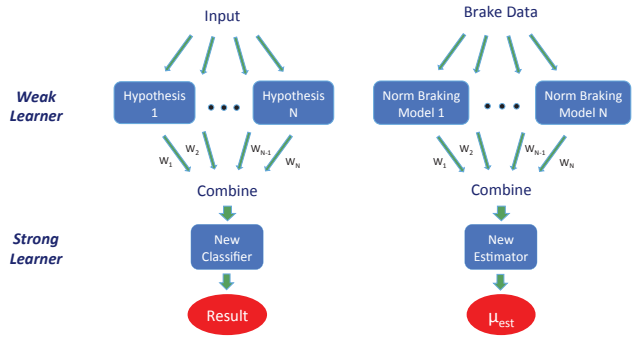

Figure 5: Ensemble learning in general (left) and in BAD (right).

even though $\Gamma_{anomaly}$ is set to a high threshold value to lower the false positive rate, if there are consecutive abnormal brake data, the score would continuously be amplified in the following steps and eventually exceed the detection threshold. Thanks to this error-amplification property of BAD, the choice of $\Gamma_{anomaly}$ is not critical.

Another portion of $e[n+1]$ is $\Delta y - \varphi[n-1]\Delta x$. That is, if there are unusually large changes in x or y within the sampling step, this can also be detected by BAD. Note that in ordinary driving, the change of slip ratio and normalized traction force would be small since the brake pedal is gradually pressed. In other words, Δx and Δy would be low.

6.2 Mid–High Slip Region

As the level of slip ratio increases and enters the mid–high slip region ($\sigma > \sigma_{linear}$), the slip-force relationship becomes non-linear. In contrast to the low slip region, brake data is dependent on the friction coefficient, thus requiring the consistency check of not only C1 but also C2. Accordingly, while BAD in the low slip region was able to be performed using a single norm braking model, BAD in the mid–high slip region cannot use the same approach.

One naive way of performing it would be as follows. First, based on the constructed norm braking models in the *norm model set*, derive the residuals of the measured brake data from each of them.[4] Next, find the norm braking model that has the minimum residual. Then, use the friction coefficient value of that norm model as the estimated friction coefficient. Finally, check whether that estimated value is consistent with the actual road condition.

However, deriving deviations from all the models in the norm model set can be a bottleneck in the detection process. Since μ is a continuous value typically ranging from 0 to 1, the cardinality of the norm model set can be infinitely large. In other words, to detect anomalies, the number of comparisons can be infinite in the worst case. This would incur high computation overhead and thus high detection delay, which is not desirable in vehicles that are typically resource-limited and require real-time detection.

Therefore, to perform BAD in real time by using only a small size of the norm model set, we use a supervised machine learning technique called *Ensemble Learning* [12]. The basic procedures for BAD in the mid–high slip region are shown in Fig. 5. Ensemble learning combines multiple base hypotheses, weak learners, in a weighted manner to form a better hypothesis, a strong learner, and predicts the status. In contrast to other supervised learning techniques such as neural networks or support vector machines (SVMs), due to its characteristic of combining weak learners, the hypotheses space can be enlarged even with a concrete finite set of base hypotheses. As a result, it incurs low computation overhead and algorithmic complexity.

[4]Due to randomness in road conditions, the brake data can conform to any of the possible norm braking models. Therefore, we must compare the data with all of them.

Algorithm 2 BAD in mid–high slip region

Inputs:
  $(x[n], y[n]) \leftarrow (\sigma[n], \rho_x[n])$
Initialize:
  $P_i[j] \leftarrow f(\phi_i[n]), j, n = 1, ..., K$
  $w_i[j] \leftarrow 1/K, j = 1, ..., K$
if $|x[n]| > \sigma_{linear}$ **then** ▷ Non-linear region
  for Each norm braking model i **do** ▷ Weak Learner
    $\phi_i[n] \leftarrow y[n] - \zeta_i(x[n])x[n]$
    $j^* \leftarrow \arg\min_j w_i[j]$
    $w_i[j^*] \leftarrow 1/K$
    $P_i[j^*] \leftarrow f(\phi_i[n])$
    $p_m \leftarrow \sum_j w_i[j]P_i[j]$
    $H_i \leftarrow \frac{1}{2}\log\frac{p_m}{1-p_m}$
    for Each training sample j **do**
      $w_i[j] \leftarrow w_i[j]\exp(-\mathrm{P_i[j]H_i})$
      normalize $w_i[j]$ so that $\sum_j w_i[j] = 1$
    end for
  end for
  if $\max H_i \approx 0$ **then** ▷ Check 1
    go to Verification stage
  else
    $\mu_{est} \leftarrow \sum(H_i\mu_i)/\sum H_i$ ▷ Strong Learner
    if $|\mu_{est} - \mu_{real}| \gg 0$ **then** ▷ Check 2
      go to Verification stage
    end if
  end if
end if

By exploiting the ensemble learning, we can achieve anomaly detection even with a small norm model set. In this paper, we construct the norm model set consisting of three different models: dry, wet, and snowy models as in Fig. 3.

Of the various ensemble learning techniques, we choose the Real AdaBoost algorithm [13] since our system deals with continuous real values. We consider each norm model in the set as a hypothesis (i.e., weak learner). In using Real AdaBoost for BAD, there are certain requirements to be met.

First, the input should be represented as a probability value. Moreover, the probability of the weak learner should be at least 0.5.

6.2.1 *Translation into a Probabilistic Value*

To satisfy the requirements, the collected brake data, (σ, ρ_x) are translated into a probability value so that it can be used as an input to the learning process. First, the residual of the measured brake data from a norm braking model i with $\mu = \mu_i$ is computed as

$$\phi_i[n] = y[n] - \zeta_i(x[n])x[n], \tag{6}$$

where

$$\zeta_i(\sigma) = \begin{cases} C_{norm} + \frac{C_{norm}^2|\sigma|}{3\mu_i} - \frac{C_{norm}^3\sigma|\sigma|}{27{\mu_i}^2} & , |\sigma| < \frac{3\mu_i}{C_{norm}} \\ \frac{\mu_i}{|\sigma|} & , \text{otherwise} \end{cases} \tag{7}$$

which is derived from Eq. (3). Here, C_{norm} is the value acquired from the norm model construction stage. The derived residual, ϕ_i, is then plugged into the probability model and translated into a probabilistic value as:

$$f(\phi_i) = \left(1 + e^{-\frac{{\phi_i}^2}{2\delta(\sigma)}}\right)/2, \quad (8)$$

which basically follows the shape of a normal distribution and has a minimum boundary of 0.5. As a result, the two requirements for Real AdaBoost input are met. The standard deviation of the probability model, δ, is set as an increasing function of the slip ratio as in $\delta(\sigma)$. This is to conform to field measurements, showing that deviations in slip-force relationship becomes larger as $|\sigma|$ increases [7].

Through translation, $f(\phi_i)$ would now represent how well the norm braking model i fits the measured brake data. In other words, the higher the value of $f(\phi_i)$ is, the more probable that the brake data was generated for a road with the friction coefficient of μ_i.

6.2.2 BAD *via Real AdaBoost*

Once the measured brake data is translated into a probabilistic form as an input for Real AdaBoost, `BAD` in the mid–high slip region is performed according to Algorithm 2.

The ensemble learning algorithm takes K training samples of brake data. For each norm braking model i, the likelihood probability $P_i[j] = f(\phi_i)$ and the hypothesis weight $w_i[j]$ for all K samples are stored in a table. Note that all brake data used in this process has to satisfy $|\sigma[n]| > \sigma_{linear}$.

Upon collecting the $(K+1)^{\text{th}}$ data, the ensemble learning process is executed. For each norm braking model i, the sample with the minimum weight in each table of i is replaced with the most recently collected data as in lines 7 and 8. This allows the learning process to take inputs of brake data with higher probability in concluding the status. Also, it allows freshness in the input data. Next, the weight of the weak learner is updated as in line 10, and the weight for each training sample is also updated. Using the derived weights of each weak learner (i.e., norm braking model), a strong learner is constructed through a weighted combination of the weak learners as in line 19. The strong learner gives an estimation of the friction coefficient, μ_{est}.

Since the method exploits a combination of weak learners, the result would be inaccurate if the diversity among them is low. As an example, three weak learners with $\mu = 0.1$, 0.2, and 0.3 will not be able to conclude on a status of $\mu = 1$. On the other hand, three weak learners with $\mu = 0.1$, 0.4, and 0.9 would have high diversity and thus be capable of representing various statuses of μ, spanning from 0 to 1.

Not only higher diversity but also more number of weak learners also enhances the accuracy of `BAD`. However, since `BAD` has a computational complexity of $\mathcal{O}(MN)$, where M is the number of weak learners and N the number of brake data samples, i.e., K, at each step, the computational complexity would also increase. Therefore, to balance between accuracy and complexity, we use three weak learners with $\mu = 0.1$, 0.4, and 0.9.

To detect anomalies, we need two checks.

- Validity of weights (line 16): At least one hypothesis in the norm model set should show a non-negligible likelihood. Failing to meet this condition means all hypotheses do not fit with the brake data. If the weak learners are carefully constructed with high diversity, such a case can only occur in case of anomalies.
- Conformance to the environment (line 20): The estimated friction coefficient from the strong learner, μ_{est}, should show a value similar to the real friction coefficient, μ_{real}. μ_{real} would reflect the actual pavement condition whereas μ_{est} would be an estimate based on the brake data. Any dissimilarity between μ_{est} and μ_{real} means that the brake data is inconsistent with the environment, thus being an anomaly.

The former requirement would be based on consistency check C1, whereas the latter would be C2. If a possible anomaly is detected by `BAD`, further verification is processed.

6.2.3 *Comparison between* μ_{est} *and* μ_{real}

Usually, μ_{est} is expected to match μ_{real}. However, for vehicle misbehavior, brake data would be abnormal and thus $\mu_{est} \neq \mu_{real}$. So, how should μ_{real} be obtained if we cannot use brake data since it might not be trustworthy?

There are many published approaches that derive μ_{real} without using them. In [14], the slip angle is derived by measuring the lateral carcass detection through a wireless piezoelectric tire sensor. By using the Brush friction model as a reference, μ_{real} is obtained through the derived slip angle. In [15], optical sensors are used to measure and analyze the spectrum of backscattered laser radiation from the road surface to obtain μ_{real}. Cameras and online road information can also be used. In [16], μ_{real} is derived by using machine learning classification. Besides camera images, online road information from a Road Weather information System (RWiS) can also help estimate the road condition [17].

Using one of these enumerated methods, μ_{real} can be obtained and used for checking the similarity to μ_{est}.

7. VERIFICATION

During the acceleration phase, tire stiffness is identified for norm model construction. Then, during the braking phase, the identified parameter is updated and used for `BAD`. The tire stiffness is dependent on inflation pressure, tire wear, temperature, and tire type. Since these factors can vary with time, there is a possibility that the tire stiffness changes between the acceleration and the braking phases. Even though the tire stiffness was accurately identified and used for norm model construction, if that value changes before entering the next braking phase, a wrong norm model would be used for detection until it is newly updated. A false-positive can thus occur during the detection phase.

Thus, when an anomaly is detected, we further need to check whether there is any significant change in the tire stiffness between the end of the acceleration phase and the start of the next braking phase.[5] To achieve this, we analyze how the dependent factors have changed. Between the two consecutive stages of acceleration and braking, tire type and wear are unlikely to change. So, we exclude them from the possible candidates to change the norm model.

On the other hand, inflation pressure and temperature may change between the two consecutive stages. Therefore, we measure whether there were any abrupt changes in the

[5]Note that any changes *during* the braking phase need not be considered since the tire stiffness is updated during that phase as well (line 20 in Algorithm 1).

two values by monitoring the measurements from the in-vehicle tire pressure monitoring sensor and in-tire temperature sensor. If the changes in those values are high, the detected anomaly is likely to be a false-positive. Otherwise, it is declared as an anomaly.

False-positive results are likely to occur due to a change in the tire stiffness during the period of transition from acceleration to braking. However, since the transition period is usually very short, such a false-positive would be unlikely.

8. EVALUATION

We now evaluate the proposed system. We first introduce the evaluation settings, and then evaluate the validity of norm model construction and the performance of BAD.

8.1 Evaluation Settings

We use CarSim [18] to obtain realistic sensor readings for the slip ratio and the normalized traction force. CarSim is high-fidelity commercial software that predicts the performance of vehicles in response to control from the driver. Due to its close agreement between simulation and field test results, it is widely used by over 30 automotive manufacturers, 60 suppliers, 150 research labs and universities. It exploits math models that characterize various aspects of the vehicles to precisely reproduce their behaviors.

We employ an E-class sedan-type car with 225/60 R18 tires and a brake system with boost and thermal effects. We adopt the following two driving scenarios, each of which is fed as an input to the CarSim simulator to obtain required sensor values.

- *Ordinary driving*: to obtain realistic values for ordinary driving, i.e., with an absolute slip ratio of up to 2%, we use the Environmental Protection Agency (EPA) Urban Dynamometer Driving Schedule (UDDS), which represents a city driving condition for 20 minutes [19].
- *Aggressive driving*: to build a driving scenario reflecting aggressive acceleration and braking of vehicles, we modify the vehicle maneuver so as to incur an absolute slip ratio of up to 30%. This scenario is used to evaluate cases with larger slip ratios.

To simulate a wide range of road conditions, we manually change the friction coefficient parameter in the simulation settings.

8.2 Validity of Norm Model Construction

Norm braking models are constructed by just measuring the brake data in the low slip region and extrapolating the data outside of the measurement region. In this subsection, we evaluate whether such extrapolated norm models also conform to the actual slip-force relationship in the mid–high slip region. To do so, we plug in the *Aggressive* driving pattern into CarSim to acquire simulated measurements in both mid–high and low slip regions, and add Gaussian noise to the simulated values to create realistic scenarios. The values for the variances of Gaussian noise were determined based on experimental data from [5, 7].

Fig. 6 shows three norm braking models constructed during the acceleration period solely by using measurements in the low slip region. *Sample*-labeled plots represent the outputs from the simulation of *Aggressive* scenario on a dry road ($\mu = 0.8$). Among the three norm models which are extrapolations with different μ values, the dry model with $\mu = 0.8$ is shown to fit well with the actual, *Sample*-labeled measurements even for high slip regions. This reveals that the proposed method of norm model construction is capable of generating the entire norm slip-force curve based only on measurements from the low slip region.

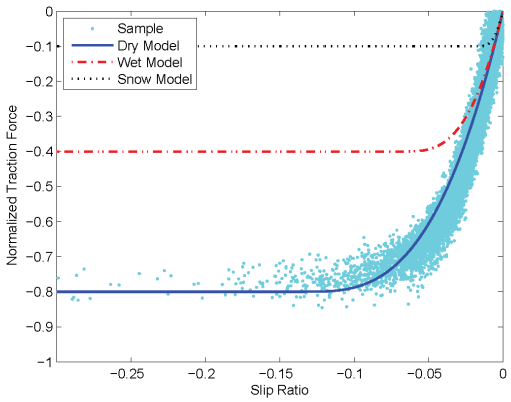

Figure 6: Conformance of the constructed norm model to the actual slip-force relationship.

8.3 Efficiency of BAD

Next, we evaluate the performance of BAD in detecting anomalies either by using identification error monitoring or ensemble learning, depending on the slip ratio value. Since the former focuses on low slip regions where $|\sigma| < \sigma_{linear}$, we use the *Ordinary* scenario. On the other hand, the latter deals with high slip regions only occurring in extreme driving cases, and hence, we use the *Aggressive* scenario for evaluation. Abnormal brake data are generated by manipulating the results obtained from simulation.

8.3.1 Detection for Low-Slip Region

In low slip regions, the identification error monitoring scheme is exploited to detect any vehicle misbehavior while braking. In Fig. 7, *Norm*-labeled data correspond to the slip-force relationship with Gaussian noise added, which indeed captures the norm braking behavior. To examine the efficiency of BAD, we inject falsified data in three different ways to yield various vehicle misbehavior cases as follows.

T1. $\sigma \approx 0$, $\rho_x < 0$. Even though the brake was applied, the slip ratio continues to have values near 0. In such case, the activation of ABS system would be bypassed when needed.

T2. $\sigma < 0$, $\rho_x \simeq 0$. The slip ratio continuously shows a negative value (i.e., instantaneous decrease of wheel speed) even though the brake is not applied.

T3. $\sigma > 0$, $\rho_x < 0$. The slip ratio shows a positive value even during the braking phase. Such a case would mean unintended acceleration.

These abnormal data were injected at 173.2, 543.5, 1178 seconds, respectively, during the braking phase in the *Ordinary* scenario. T1 is set to quickly variate from the true data, whereas T2 and T3 are set to slowly drift away from it. A forgetting factor of 0.9994 was used for the RLS operation.

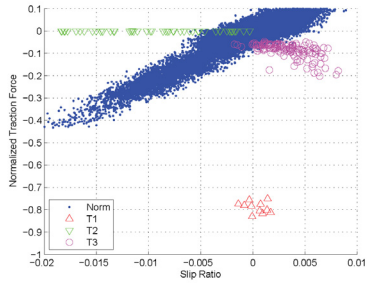

Figure 7: Brake data used for evaluating the performance of BAD in low slip region.

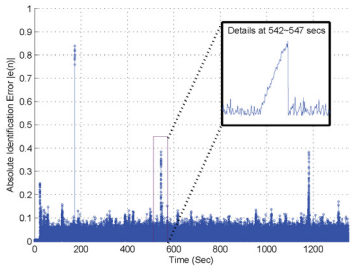

Figure 8: Determined anomaly scores from BAD in low slip region.

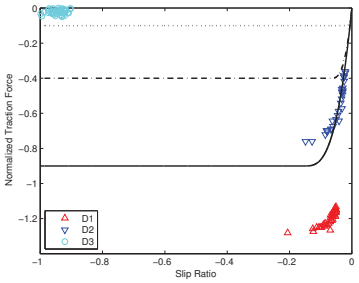

Figure 9: Brake data used for evaluating the performance of BAD in mid–high slip region.

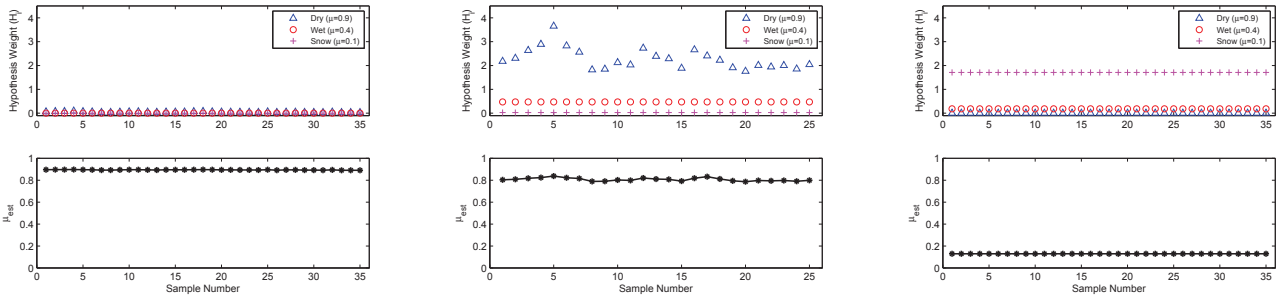

(a) Results for D1 (b) Results for D2 (c) Results for D3

Figure 10: Hypotheses' weights of weak learners and the output from the strong learner in BAD.

The results of running BAD in the low slip region are plotted in Fig. 8, showing three notable peaks of anomaly scores. Each of them occurs at the point where the falsified brake data were inserted. Thus, by using BAD, the listed vehicle anomalies can be detected with $\Gamma_{anomaly} = 0.3$ (or a higher value if the duration of misbehavior was increased). The initial anomaly score is slightly high due to insufficient data for learning the tire stiffness, thus not being able to converge. However, the convergence delay can be adjusted from the RLS parameter $P[0] = \delta I$.

Fig. 8 also shows the magnified plot of anomaly scores during 542∼547 seconds. It can be seen that consecutive vehicle misbehavior leads to continuous increase in the anomaly score, conforming to the analysis on the error-amplification property of BAD.

8.3.2 *Detection for Mid–High Slip Region*

In mid–high slip regions, ensemble learning is used to learn the weights of each norm model and construct a strong learner that concludes on the estimated friction coefficient. Verification of the maximum hypothesis weight and comparison between the estimated and actual friction coefficient are performed. To evaluate BAD in the mid–high slip region, the *Aggressive* driving scenario was simulated using CarSim, part of which was manipulated to emulate the vehicle misbehavior. The table size of weak learner was set as $K = 5$, and three weak learners with $\mu = 0.1$, 0.4, and 0.9 were used. Fig. 9 shows the data points that were examined. Each data point lies in the mid–high slip region and represents the following.

D1. Brake data that deviates significantly from all norm models. This can happen if the vehicle does not get enough slip for deceleration, even though the brake is pressed hard.

D2. Brake data from a vehicle running on a dry road of μ=0.8. This would represent the case where the data does not exactly conform to any of the norm models but still has a valid friction coefficient.

D3. Even though the brake pedal is pressed little, the slip ratio tends to have a value near -1, which in turn might activate the ABS.

The anomaly detection results for the dataset are shown in Figures 10(a) to 10(c), where the upper subplot shows the hypothesis weights, i.e., weights of the weak learners, and the lower subplot shows the estimated friction coefficient for each brake data, i.e., the output from the strong learner. Moreover, the results are summarized in Table 1.

Test No.	H_i validity	μ_{est}	Result
D1	Invalid	0.89	C1 Fail
D2	Valid	0.79	C2 Check ($\mu_{est} \approx \mu_{real}$?)
D3	Valid	0.14	C2 Check ($\mu_{est} \approx \mu_{real}$?)

Table 1: Summary of BAD in mid–high slip region

First, for the results from D1 shown in Fig. 10(a), although μ_{est} (≈ 0.9) seems valid, all the hypothesis weights show a value near 0, meaning invalid hypothesis weights. In other words, such data would be determined as anomaly since Check 1 (line 16 of Algorithm 2) would fail.

Next, results for D2 and D3 are shown in Figures 10(b) and 10(c), respectively. The system is shown to accurately estimate the friction coefficient using the ensemble learning. Since the maximum hypothesis weights, $\max H_i$, for both cases are non-negligible (i.e., valid hypothesis weights), Check 1 will pass. Thus, Check 2 (line 20 of Algorithm 2) should be performed further. Suppose D2 occurred while driving on a snowy road yielding low μ_{real}, and D3 occurred while driving on a dry asphalt road yielding high μ_{real}. Then, through Check 2, since $\mu_{est} \neq \mu_{real}$, the brake data would be declared as an anomaly.

9. CONCLUSION

Advanced vehicle technologies have brought new threats and dangers such as sudden unintended acceleration or braking due to software bugs or hardware glitches. In this paper, after constructing a norm braking model, we proposed a real-time anomaly detection method, called `BAD`, which compares the brake data with the norm model to detect any anomalies in the Brake-by-Wire system. `BAD` exploits both RLS error monitoring in low slip regions and ensemble learning in mid–high slip regions, and is supplemented by a verification process to further check the anomalies, thereby minimizing false-positives. Our evaluation results demonstrated that the proposed methods are indeed valid and efficient. In the future, we would like to analyze root causes and conduct field measurements. Also, it would be interesting to explore the recovery actions to be taken upon detection of anomalies.

Acknowledgements

The work reported in this paper was supported in part by the US NSF under Grants CNS1329702 and CNS1446117, and the DGIST R&D Program of MSIP of Korea (CPS Global Center) and the Global Research Laboratory Program through NRF funded by MSIP of Korea (2013K1A1A2 A02078326).

10. REFERENCES

[1] Barr's trial testimony - Safety Research & Strategies, http://www.safetyresearch.net/
[2] R. Belt, "An Electronic Cause for Sudden Unintended Acceleration", http://www.autosafety.org
[3] H. Pacejka and E. Bakker, "The Magic Formula Tyre Model," in *Vehicle System Dynamics: International Journal of Vehicle Mechanics and Mobility*, 1992.
[4] E. Bakker, L. Nyborg, and H. Pacejka, "Tyre modelling for use in vehicle dynamics studies," in *SAE Techinical paper 870421*, Feb. 1987.
[5] S. Muller, M. Uchanski, and K. Hedrick, "Estimation of the Maximum Tire-Road Friction Coefficient," in *ASME Journal of Dynamic Systems, Measurement, and Control*, Dec. 2003.
[6] J. Hahn, R. Rajamani, and L. Alexander, "GPS-Based Real-Time Identification of Tire-Road Friction Coefficient," in *IEEE Transactions on Control Systems Technology*, May 2002.
[7] J. Svendenius, "Tire modeling and friction estimation," Ph.D. thesis, Department Automatic Control, Lund University. 2007.
[8] R. Rajamani, G. Phanomchoeng, D. Piyabongkarn, and J. Lew, "Algorithms for Real-Time Estimation of Individual Wheel Tire-Road Friction Coefficients," in *IEEE Transactions on Mechatronics,"* in Dec. 2012.
[9] R. Rajamani, *Vehicle Dynamics and Control.* New York: Springer-Verlag, 2005.
[10] C. Carlson and J. Gerdes, "Consistent Nonlinear Estimation of Longitudinal Tire Stiffness and Effective Radius," in *IEEE Transactions on Control Systems Technology*, Nov. 2005.
[11] S. Haykin, Adaptive Filter Theory, 2nd ed. Prentice-Hall, 1991.
[12] T. Dietterich, "Ensemble methods in machine learning," in *Multiple Classifier Systems*, 2000.
[13] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting," in *Annals of statistics*, 2000.
[14] G. Erdogan, L. Alexander, and R. Rajamani, "Estimation of Tire-Road Friction Coefficient Using a Novel Wireless Piezoelectric Tire Sensor," in *IEEE Sensors Journal,* 2011.
[15] J. Casselgren, M. Sjodahl, S. Woxneryd, and M. Sanfridsson, "Classification of Road Conditions - to Improve Safety," in *Advanced Microsystems for Automotive Applications,* 2007.
[16] P. Jonsson, "Classification of road conditions: From camera images and weather data," in *IEEE CIMSA,* 2011.
[17] http://ops.fhwa.dot.gov/weather/
[18] http://www.carsim.com/
[19] http://www.epa.gov/nvfel/testing/dynamometer.htm

APPENDIX

A. SLIP RATIO DERIVATION

The slip ratio, $x[n] = \sigma$, is derived from the measured wheel speed, $r_{eff}w_w$, and the vehicle speed, V. The wheel angular velocity, w_w, is measured by the wheel speed sensors that monitor the time-varying magnetic flux of the tone wheel and convert it to angular velocity. The effective rolling tire radius, r_{eff}, can be easily measured as in [10] or assumed to be given [8]. Finally, the vehicle speed can be measured using the non-driven wheel speed sensors or an accelerometer with assistance from a carrier-phase-based GPS system [8]. Such measurements are broadcast on the Controller Area Network (CAN) bus for various vehicle operations, and thus can be easily acquired.

B. NORMALIZED TRACTION FORCE DERIVATION

To derive the normalized traction force, $y[n] = \rho_x$, the longitudinal and the vertical normal forces have to be calculated or measured. The longitudinal force at each tire, $F_{x,i}$, can be calculated based on the rotational dynamics of each wheel as

$$F_{x,i} = \frac{T_{engine,i} - T_{brake,i} - I_w \dot{w}_{w,i}}{r_{eff}}, \tag{9}$$

where the subscripts $i = fl, fr, rl, rr$ represent four wheels, $T_{engine,i}$ denotes the engine torque, $T_{brake,i}$ the brake torque, and I_w the inertia of the wheel. Here, I_w, r_{eff}, and $w_{w,i}$ are variables that can be easily measured or assumed to be given [8]. The engine and the brake torque values are typically available over the CAN bus; if not, they can be derived by measuring the accelerator and brake pedal position and converting them to torque values using the driver demand torque map stored in each subsystem.

To derive the vertical normal force at each tire, the static force model of the vehicle can be used as in [9]:

$$F_{z,fl} = F_{z,fr} = \frac{mgL_r - mh\dot{V} - C_a V^2 h_a}{2L}, \tag{10}$$

$$F_{z,rl} = F_{z,rr} = \frac{mgL_f + mh\dot{V} + C_a V^2 h_a}{2L}, \tag{11}$$

where m denotes the vehicle mass, C_a the aerodynamic drag parameter, g the acceleration of gravity, and L_f, L_r, the distance from the front and rear tire, respectively. Moreover, h and h_a denote the height of the center of gravity and of the aerodynamic drag force. Each variable on the right-hand side can be measured, calculated, and pre-determined [8,9]. Depending on the available sensors, the normal force can also be derived by exploiting the vehicle pitch dynamics [5] or wheel suspension dynamics [6]. The details of these are omitted due to space limitation.

In conclusion, using various sensor measurements available via the CAN bus, the normalized traction force can be determined. The results in [5,6,8,9] show that the numerically derived values match well with the actual ones.