**QUEEN'S
UNIVERSITY
BELFAST**

Reliability Validation of Learning Enabled Vehicle Tracking

Sun, Y., Zhou, Y., Maskell, S., Sharp, J., & Huang, X. (Accepted/In press). *Reliability Validation of Learning Enabled Vehicle Tracking*. Paper presented at International Conference on Robotics and Automation, Paris, France.

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights
Copyright 2020 IEEE. This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights
Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Reliability Validation of Learning Enabled Vehicle Tracking

Youcheng Sun^{1*}, Yifan Zhou^{2*}, Simon Maskell², James Sharp³ and Xiaowei Huang²

Abstract—This paper studies the reliability of a real-world learning-enabled system, which conducts dynamic vehicle tracking based on a high-resolution wide-area motion imagery input. The system consists of multiple neural network components – to process the imagery inputs – and multiple symbolic (Kalman filter) components – to analyse the processed information for vehicle tracking. It is known that neural networks suffer from adversarial examples, which make them lack robustness. However, it is unclear if and how the adversarial examples over learning components can affect the overall system-level reliability. By integrating a coverage-guided neural network testing tool, DeepConcolic, with the vehicle tracking system, we found that (1) the overall system can be resilient to some adversarial examples thanks to the existence of other components, and (2) the overall system presents an extra level of uncertainty which cannot be determined by analysing the deep learning components only. This research suggests the need for novel verification and validation methods for learning-enabled systems.

I. INTRODUCTION

The wide-scale deployment of Deep Neural Networks (DNNs) in safety-critical applications, such as self-driving cars, healthcare and Unmanned Air Vehicles (UAVs), increases the demand for tools that can test, validate, verify, and ultimately certify such systems [1]. Normally, autonomous systems, or more specifically learning-enabled systems, contain both data-driven learning components and model-based non-learning components [2], and a certification needs to consider both categories of components, and the system as a whole.

Structural code coverage combined with requirements coverage has been the primary approach for measuring test completeness as assurance evidence to support safety arguments, for the certification of safety-critical software. Testing techniques for machine learning components, primarily DNNs, are comparatively new and have only been actively developed in the past few years, e.g., [3], [4]. Unlike model-based software systems, DNNs are usually considered as black boxes, and therefore it is difficult to understand their behaviour by means of inspection. In [5], we developed a tool, DeepConcolic, to work with a number of extensions to the MC/DC coverage metric, targeting DNNs. The MC/DC coverage metric [6] is recommended in a number of certification documents and standards, such

as RTCA's DO-178C and ISO26262, for the development of safety critical software. Its extension to DNN testing has been shown to be successful in testing DNNs by utilising white-box testing methods, i.e., by exercising the known structure with the parameters of a DNN to gather assurance evidence. It is, however, unclear whether such a testing tool can still be effective when working with learning-enabled systems containing both learning and model-based components. Primarily, we want to understand the following two research questions:

- Q1 Can the system as a whole be resilient against the deficits discovered over the learning components?
- Q2 Is there new uncertainty needed to be considered in terms of the interaction between learning and non-learning components?

The key motivation for considering Q1 is to understand whether, when generating a test suite, a DNN testing tool needs to consider the existence of other components in order to assess the safety of the system. The key motivation for considering Q2 is to understand whether a DNN testing tool can take advantage of the uncertainty presented in the interaction between learning and non-learning components in generating test cases.

Specifically, for the learning-enabled systems, we consider in this paper a tracking system in Wide Area Motion Imagery (WAMI) [7], where the vehicle detection is implemented by a few DNNs and the tracking is implemented with a Kalman filter based method. Using this system, we consider its reliability when running in an adversarial environment, where an adversary can have limited access to the system by intercepting the inputs to the perception unit.

Our experiments provide affirmative answers to the above research questions and point out the urgent need to develop system-level testing tools to support the certification of learning-enabled systems.

II. PRELIMINARIES

A (feedforward and deep) neural network $\mathcal{N}$ is a function that maps an input x to an output. According to the tasks, the output can be of different format. For example, for classification task, the DNN computes a *label*, which is denoted by $\mathcal{N}(x).label$.

Adversarial examples [8] are two very similar inputs with different labels. The existence of such pairs has been used as a proxy metric for the training quality of a DNN. Given an input x_1 that is correctly labeled by a DNN $\mathcal{N}$, another input x_2 is said to be an adversarial example if x_1 and x_2 are “close enough”, i.e., $\|x_1 - x_2\|_p \leq \epsilon$, and $\mathcal{N}(x_1).label \neq$

Youcheng Sun and Yifan Zhou contribute equally to this work. This work is also partially supported by the UK EPSRC through the Offshore Robotics for Certification of Assets (ORCA) [EP/R026173/1]

¹ School of Electrical Engineering, Electronics and Computer Science, Queen's University Belfast, UK

² School of Electrical Engineering, Electronics and Computer Science, University of Liverpool, UK

³ Defence Science and Technology Laboratory (Dstl), UK

$\mathcal{N}(x_2).label$. Here, p denotes the L^p -norm distance metric and ϵ measures the term “sufficiently small”.

A number of algorithms have been proposed to find adversarial examples. However, such methods are not able to quantify progress, and thus, are not useful as a stopping condition for testing a DNN. This has motivated the coverage criteria recently developed for DNNs.

A. The DeepConcolic Tool

Several structural coverage criteria have been designed for DNNs, including neuron coverage [3] and its extensions [9], and variants of Multiple Condition/Decision Coverage (MC/DC) for DNNs [10]. These coverage criteria quantify the exhaustiveness of a test suite for a DNN. Neuron coverage requires that, for every neuron in the network, there must exist at least one test case in the generated test suite that lifts its activation value above some threshold; the criteria in [9] generalise the neuron coverage from a single neuron to a set of neurons. The MC/DC variants for DNNs capture the fact that a (decision) feature (a set of neurons) in a layer is directly decided by its connected (condition) features in the previous layer, and thus require test conditions such that every condition feature must be exhibited regardless of its effect on the decision feature.

DeepConcolic¹ implements a concolic analysis that examines the set of behaviors of a DNN, and is able to identify potentially problematic input/output pairs. DeepConcolic generates test cases and adversarial examples for DNNs following the specified test conditions. Developers can use the test results to compare different DNN models, and the adversarial examples can be used to improve and re-train the DNN or develop an adversarial example mitigation strategy. Moreover, a major safety challenge for the use of DNNs is due to the lack of understanding about how a decision is made by a DNN. The DeepConcolic tool is able to test each particular component of a DNN, and this aids human analysis of the internal structures of a DNN. By understanding these structures, we improve the confidence in DNN behaviour; this is different from providing guarantees (for example [11], [12], [13], [14], [15]), just as testing conventional software does not provide guarantees of correctness.

III. DETECTION FROM WIDE-AREA MOTION IMAGERY

This section describes the technical details of a tracking system. The tracking system requires continuous imagery input from e.g., airborne high-resolution cameras. The input to the tracking system is a video, which consists of a finite sequence of images. Each image contains a number of vehicles. Similar to [7], we use the WPAFB 2009 [16] dataset. The images were taken by a camera system with six optical sensors that had already been stitched to cover a wide area of around 35 km^2 . The frame rate is 1.25Hz. This dataset includes 1025 frames, which is around 13 minutes of video. It is divided into training video (512 frames) and testing video (513 frames). All the vehicles and their trajectories

are manually annotated. There are multiple resolutions of videos in the dataset. For the experiment, we chose to use the $12,000 \times 10,000$ images, in which the size of vehicles is smaller than 10×10 pixels. We use α_i to denote the i -th frame and $\alpha_i(x, y)$ the pixel on the intersection of x -th column and y -th row of α_i .

In the following, we explain how the tracking system works by having a video as input. In general, this is done in two steps: detection and tracking. In Section III-A through to Section III-D, we explain the detection steps, i.e., how to detect a vehicle with CNN-based perception units. This is followed by the tracking step in Section III-E.

A. Background Construction

Vehicle detection in WAMI video is a challenging task due to the lack of vehicle appearances and the existence of frequent pixel noises. It has been discussed in [17], [18] that an *appearance-based* object detector may cause a large number of false alarms. For this reason, in this paper, we only consider *detecting moving objects* for tracking.

Background construction is a fundamental step in extracting pixel changes from the input image. The background is built for the current environment from a number of previous frames that were captured by a moving camera system. It proceeds in the following steps.

a) *Image registration*: is to compensate for the camera motion by aligning all the previous frames to the current frame. The key is to estimate a transformation matrix, h_t^{t-k} , which transforms frame α_{t-k} to frame α_t using a given transformation function. For the transformation function, we consider projective transformation (or homography), which has been widely applied in multi-perspective geometry, an area where WAMI camera systems are already utilised.

The estimation of h_t^{t-k} is generated by applying feature-based approaches. First of all, feature points from images at frame α_{t-k} and α_t , respectively, are extracted by feature detectors (e.g., Harris corner or SIFT-like [19] approaches). Second, feature descriptors, such as SURF [20] and ORB [21] are computed for all detected feature points. Finally, pairs of corresponding feature points between two images can be identified and the matrix h_t^{t-k} can be estimated by using RANSAC [22] which is robust against outliers.

b) *Background Modeling*: We generate the background, I_t^{bg} , for each time t , by computing the median image of the L previously-aligned frames, i.e.,

$$I_t^{bg}(x, y) = \frac{1}{L} \sum_{i=1}^L \alpha_{t-i}(x, y) \quad (1)$$

In our experiments, we take either $L = 4$ or $L = 5$.

Note that, to align the L previous frames to the newly received frame, only one image registration process is performed. After obtaining the matrices $h_{t-1}^{t-2}, h_{t-1}^{t-3}, \dots$ by processing previous frames, we perform image registration once to get h_t^{t-1} , and then let

$$h_t^{t-2} = h_t^{t-1} \times h_{t-1}^{t-2}, \quad h_t^{t-3} = h_t^{t-1} \times h_{t-1}^{t-3}. \quad (2)$$

¹<https://github.com/TrustAI/DeepConcolic>

c) *Extraction of Potential Moving Objects*: By comparing the difference between I_t^{bg} and the current frame α_t , we can extract a set Q_{bc} of potential moving objects by first computing the following set of pixels

$$P_{bc} = \{(x, y) \mid |I_t^{bg}(x, y) - \alpha_t(x, y)| > \delta_{bc}, (x, y) \in \Gamma\} \quad (3)$$

and then applying image morphology operation on P_{bc} , where Γ is the set of pixels and δ_{bc} is a threshold value to determine which pixels should be considered.

B. CNN for Detection Refinement

After obtaining P_{bc} , we develop a Convolutional Neural Network (CNN) $\mathcal{N}_{dr}$, a type of DNN, to detect vehicles. We highlight a few design decisions. The major causes of false alarms generated by the background subtraction are: poor image registration, light changes and the apparent displacement of high objects (e.g., buildings and trees) caused by parallax. We emphasise that the objects of interest (e.g., vehicles) mostly, but not exclusively, appear on roads. Moreover, we perceive that a moving object generates a temporal pattern (e.g., a track) that can be exploited to discern whether or not a detection is an object of interest. Thus, in addition to the shape of the vehicle in the current frame, we assert that the historical context of the same place can help to distinguish the objects of interest and false alarms.

By the above observations, we create a binary classification CNN $\mathcal{N}_{dr} : \mathbb{R}^{21 \times 21 \times (N+1)} \rightarrow \{0, 1\}$ to predict whether a 21×21 pixels window contains a moving object given aligned image patches generated from the previous N frames. The 21×21 pixels window is identified by considering the image patches from the set Q_{bc} . We suggest $N = 3$ in this paper, as it is the maximum time for a vehicle to cross the window. The input to the CNN is a $21 \times 21 \times (N + 1)$ matrix and the convolutional layers are identical to the traditional 2D CNNs except that the three colour channels are substituted with $N + 1$ grey-level frames.

Essentially, $\mathcal{N}_{dr}$ acts as a filter to remove from Q_{bc} objects that are unlikely to be vehicles. Let Q_{dr} be the obtained set of moving objects. If the size of an image patch in Q_{dr} is similar to a vehicle, we directly label it as a vehicle. On the other hand, if the size of the image patch in Q_{dr} is larger than a vehicle, i.e., there may be multiple vehicles, we pass this image patch to the location prediction for further processing.

C. CNN for Location Prediction

We take a regression CNN $\mathcal{N}_{lp} : \mathbb{R}^{45 \times 45 \times (N+1)} \rightarrow \mathbb{R}^{15 \times 15}$ to process image patches passed over from the detection refinement phase. As in [18], a regression CNN can predict the locations of objects given spatial and temporal information. The input to $\mathcal{N}_{lp}$ is similar to the classification CNN $\mathcal{N}_{dr}$ described in Section III-B, except that the size of the window is enlarged to 45×45 . The output of $\mathcal{N}_{lp}$ is a 225-dimensional vector, equivalent to a down-sampled image (15×15) for reducing computational cost.

For each 15×15 image, we apply a filter to obtain those pixels whose values are greater than not only a threshold value δ_{lp} but also the values of its adjacent pixels. We then

obtain another 15×15 image with a few bright pixels, each of which is labelled as a vehicle. Let O be the set of moving objects updated from Q_{dr} after applying location prediction.

D. Detection Framework

The processing chain of the detector is shown in Figure 1(a). At the beginning of the video, the detector takes the first L frames to construct the background, thus the detections from frame $L + 1$ can be generated. After the detection process finishes in each iteration, it is added to the template of previous frames. The updating process substitutes the oldest frame with the input frame. This is to ensure that the background always considers the latest scene, since the frame rate is usually low in WAMI videos such that parallax effects and light changes can be pronounced. As we wish to detect very small and unclear vehicles, we apply a small background subtraction threshold and a minimum blob size. This, therefore, leads to a huge number of potential blobs. The classification CNN is used to accept a number of blobs. As mentioned in Section III-B, the CNN only predicts if the window contains a moving object or not. According to our experiments, the cases where multiple blobs belong to one vehicle and one blob includes multiple vehicles, occur frequently. Thus, we design two corresponding scenarios: the blob is very close to another blob(s); the size of the blob is larger than 20×20 . If any blob follows either of the two scenarios, we do not consider the blob for output. The regression CNN (Section III-C) is performed on these scenarios to predict the locations of the vehicles in the corresponding region, and a default blob will be given. If the blob does not follow any of the scenarios, this blob will be outputted directly as a detection. Finally, the detected vehicles include the output of both sets.

E. Object Tracking

1) *Problem Statement*: We consider a single target tracker (ie. Kalman filter) to track a vehicle given all the detection points over time in the field of view. The track is initialised by manually giving a starting point and a zero initial velocity, such that the state vector is defined as $s_t = [x_t, y_t, 0, 0]^T$ where $[x_t, y_t]$ is the coordinate of the starting point. We define the initial covariance of the target, $P = \text{diag}[30, 30, 20, 20]^2$, which is the initial uncertainty of the target's state².

A near-constant velocity model is applied as the dynamic model in the Kalman filter which is defined in (4).

$$s_{t|t-1} = F \cdot s_{t-1} + \omega_t \quad \text{s.t. } F = \begin{bmatrix} \mathbb{I}_{2 \times 2} & \mathbb{I}_{2 \times 2} \\ \mathbb{O}_{2 \times 2} & \mathbb{I}_{2 \times 2} \end{bmatrix} \quad (4)$$

where $\mathbb{I}$ is a identity matrix, $\mathbb{O}$ is a zero matrix, s_{t-1} is the state vector in previous timestep, $s_{t|t-1}$ is the predicted state vector and ω_t is the process noise which can be further defined as $\omega_t \sim \mathcal{N}(0, Q)$, where $\mathcal{N}(0, Q)$ denotes a Gaussian

²With this configuration it is not necessary for the starting point to be a precise position of a vehicle, and the tracker will find a proximate target to track on. However, it is possible to define a specific velocity and reduce the uncertainty in P , so the tracker can track a particular target.

Fig. 1: (a) The architecture of the vehicle detector. (b) Workflow for testing the WAMI tracking system.

distribution whose mean is zero and the covariance is Q defined in (5).

$$Q = \sigma_q^2 \cdot \begin{bmatrix} \frac{1}{3}dt^3 \cdot \mathbb{I}_{2 \times 2} & \frac{1}{2}dt^2 \cdot \mathbb{I}_{2 \times 2} \\ \frac{1}{2}dt^2 \cdot \mathbb{I}_{2 \times 2} & \mathbb{I}_{2 \times 2} \end{bmatrix} \quad (5)$$

where dt is the time interval between two frames and σ_q is a configurable constant. $\sigma_q = 3$ is suggested for the aforementioned WAMI video.

Next, we define the measurement model as (6).

$$z_t = H \cdot s_t + \epsilon_t \quad \text{s.t.} \quad H = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad (6)$$

where z_t is the measurement (which is the position of the tracked vehicle), s_t denotes the true state of the vehicle and $\epsilon_t \sim \mathcal{N}(0, R)$ denotes the measurement noise which models the uncertainty involved in the detection process. R is defined as $R = \sigma_r^2 \cdot \mathbb{I}_{2 \times 2}$, where we suggest $\sigma_r = 5$ for the WAMI video.

Since the camera system is moving, the position should be compensated for such motion using the identical transformation function for image registration. However, we ignore the influence to the velocity as it is relatively small, but consider integrating this into the process noise.

2) *Measurement Association*: During the update step of the Kalman filter, the residual measurement should be calculated by subtracting the measurement (z_t) from the predicted state ($s_{t|t-1}$). In the tracking system, the gated nearest neighbour to obtain the measurement from a set of detections is considered. K-nearest neighbour is firstly applied to find the nearest detection, $\hat{z}_t$, of the predicted measurement, $H \cdot s_{t|t-1}$. Then the Mahalanobis distance between $\hat{z}_t$ and $H \cdot s_{t|t-1}$ is calculated as follows:

$$D_t = \sqrt{(\hat{z}_t - H \cdot s_{t|t-1})^T \cdot P_{t|t-1}^{-1} \cdot (\hat{z}_t - H \cdot s_{t|t-1})} \quad (7)$$

where $P_{t|t-1}$ is the Innovation covariance, which is defined within Kalman filter.

A potential measurement is adopted if $D_t \leq g$ with $g = 2$ in our experiment. If there is no available measurement, the update step will not be performed and the state uncertainty accumulates. It can be noticed that a large covariance leads to a large search window. Because the search window can be unreasonably large, we halt the tracking process when the trace of the covariance matrix exceeds a pre-determined value.

IV. RELIABILITY TESTING FRAMEWORK

We consider the reliability of the vehicle tracking system introduced above when its perception units are subject to adversarial attacks. In general, we assume that the vehicle tracking system is running in an adversarial environment and the adversary is able to intercept the inputs to the perception unit in a limited way.

Figure 1(b) outlines the workflow of our testing framework, where DeepConcolic is deployed for reliability testing of the vehicle tracking system. Inside the dashed block, resides the workflow of the original WAMI tracking system as described in Section III. In order to test the reliability of this vehicle tracker, we interface it with DeepConcolic, which accepts as inputs the original WAMI image inputs and the convolutional network. It then generates the distortion, via MC/DC testing, for the input image to lead the network into failing to detect vehicles.

A. Tracks

First of all, we formalise the concept of a track implemented in the WAMI tracker. At each step (frame), every detected vehicle can be identified with a *location*, represented as a tuple, $l = (x, y)$, where x and y are the location's horizontal and vertical coordinates.

Given two locations $l_1 = (x_1, y_1)$ and $l_2 = (x_2, y_2)$, we assume that there is a distance function $dist_{loc}(l_1, l_2)$ to quantify the distance between the two detected locations. For example, we can define the distance function as follows.

$$dist_{loc}(l_1, l_2) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$$

In reality, when tracking a car, it is reasonable to assume a threshold ϵ such that as long as the distance between two locations does not exceed ϵ , i.e., $dist_{loc}(l_1, l_2) \leq \epsilon$, the two locations l_1 and l_2 can be regarded as the same.

A detected *track*, denoted by $\vec{r}$, consists of a sequence of locations such that $\vec{r} = l_1 \dots l_n$, where n is the total number of steps for tracking. We write the number of steps $|\vec{r}|$ as the length of the track. Let $\mathcal{R}$ be the set of finite tracks. Subsequently, given two tracks $\vec{r}_1$ and $\vec{r}_2$ of the same length, there are multiple choices to measure the difference between the two, denoted by $dist_{tr}(\vec{r}_1, \vec{r}_2)$, including

$$\begin{aligned} dist_{tr}^{sum}(\vec{r}_1, \vec{r}_2) &= \sum_{1 \leq i \leq n} dist_{loc}(l_{1,i}, l_{2,i}) \\ dist_{tr}^{mean}(\vec{r}_1, \vec{r}_2) &= \frac{1}{n} \cdot dist_{tr}^{sum}(\vec{r}_1, \vec{r}_2) \\ dist_{tr}^{max}(\vec{r}_1, \vec{r}_2) &= \max_{1 \leq i \leq n} \{dist_{loc}(l_{1,i}, l_{2,i})\} \\ dist_{tr}^{\epsilon}(\vec{r}_1, \vec{r}_2) &= |\{dist_{loc}(l_{1,i}, l_{2,i}) > \epsilon | 1 \leq i \leq n\}| \end{aligned}$$

Fig. 2: Illustration of the dependencies and workflows between data, for the cases of before robustness testing (Left) and after robustness testing (Right), respectively.

where $l_{m,i}$ represents the i -th location on the track $\vec{r}_m$. Intuitively, $dist_{tr}^{sum}(\vec{r}_1, \vec{r}_2)$ expresses the accumulated distance between locations of two tracks, $dist_{tr}^{mean}(\vec{r}_1, \vec{r}_2)$ is the average distance between locations of two tracks, $dist_{tr}^{max}(\vec{r}_1, \vec{r}_2)$ highlights the largest distance between locations of two tracks, and $dist_{tr}^{\epsilon}(\vec{r}_1, \vec{r}_2)$ counts the number of steps on which the distance between locations of two tracks is greater than ϵ .

B. Reliability Definition

We start from explaining how to obtain a track by the WAMI detector in Section III. Given a number n , we let X^n be the set of image sequences of length n . For each image sequence $\alpha_1 \dots \alpha_n \in X^n$, the WAMI tracking system maps it into a track as follows. Let α_{k+1} be the $(k+1)$ -th image and $s_k = (l_k, C_k)$ the state of the Kalman filter after k steps, where l_k is the location, and C_k is the covariance matrix representing the uncertainty of l_k . From s_k , an estimated location $\hat{l}_{k+1}$ can be obtained. Moreover, from the image α_{k+1} , by the WAMI detection based on both $\mathcal{N}_{dr}$ and $\mathcal{N}_{lp}$ in Section III, we can have a set of locations Q_{k+1} representing the detected moving objects. By selecting a location $q_{k+1} \in Q_{k+1}$ which is closest to $\hat{l}_{k+1}$, we determine the next location l_{k+1} , and use this information to update C_k into C_{k+1} . In the left of Figure 2, the causality between the above entities is exhibited. For simplicity, we can write the resulting track as

$$\vec{r} = G(\alpha_1 \dots \alpha_n) = l_1 \dots l_n \quad (8)$$

Because of the existence of the adversary, when the detection components $\mathcal{N}_{dr}$ and $\mathcal{N}_{lp}$ lack robustness, from α_{k+1} we may have a different set, Q'_{k+1} , of detected moving objects. Therefore, there may be a different $q'_{k+1} \neq q_{k+1}$, which will in turn lead to different l'_{k+1} and C'_{k+1} . The causality of these entities is exhibited in the right of Figure 2. We write the resulting track as

$$\vec{r}' = (adv \odot G)(\alpha_1 \dots \alpha_n) = l'_1 \dots l'_n \quad (9)$$

Therefore, given a track $\vec{r}$, by attacking the learning components of the tracking system, another track $\vec{r}'$ may be returned. Now, we can define the reliability of a tracking system. The reliability of a tracking system over an image sequence, $\alpha_1 \dots \alpha_n$, a function, $dist_{tr}$, and a threshold, E , is defined as the non-existence of an adversary, adv , such that

$$dist_{tr}(G(\alpha_1 \dots \alpha_n), (adv \odot G)(\alpha_1 \dots \alpha_n)) > E. \quad (10)$$

when $\|\alpha_1 \dots \alpha_n - adv(\alpha_1 \dots \alpha_n)\|_p \leq \epsilon_0$, with $\epsilon_0 > 0$ being a small enough real number. We use $\|\alpha_1 \dots \alpha_n - adv(\alpha_1 \dots \alpha_n)\|_p \leq \epsilon_0$ to denote that the maximum adversarial

perturbation to be considered is no more than ϵ_0 , when measured with L^p -norm. Under such a constraint, the reliability requirement is to ensure that no adversary can produce an attack that results in the attacked track significantly deviating from the original track, i.e., $dist_{tr}(G(\alpha_1 \dots \alpha_n), (adv \odot G)(\alpha_1 \dots \alpha_n)) > E$.

C. Reliability Validation

We use DeepConcolic generated test cases for the detection components $\mathcal{N}_{dr}$ and $\mathcal{N}_{lp}$ to validate the reliability of the tracking system. In theory, we may generate test cases for all steps $k > 1$ or a few randomly selected steps. In practice, we found that the tracking system in Figure 1 (b) is more sensitive to the scheme when there are consecutive missing detections of the vehicle.

Thus, we design the testing scheme $Test(s, k)$ to test the tracking system, that is, starting from the s -th frame, test cases are generated by applying DeepConcolic to k successive frames for the tracking system. A track comprises of a finite sequence of vehicle detection, and $Test(s, k)$ is useful to access the reliability of different parts of the track. This helps reveal the potential vulnerabilities of the tracking system.

V. EXPERIMENTS

We sample a number of tracks with maximum length of 30, and then apply the $Test(s, k)$ with a variety of configurations: $s \in [5, 20]$, $k \in [1, 4]$. Example tracks are as in Figure 3.

Fig. 3: Original detected tracks from the tracking system

At first, we confirm that even by only testing the deep learning component in the tracking system, it is able to find test inputs that help discover the vulnerabilities in the tracking. Figure 4 shows the adversarial tracks found by the testing.

Furthermore, different parts of the same track exhibit different levels of robustness. For example, as shown in Figure 5, the frames x_5, x_6 are less robust than x_{11}, x_{12} , as

Fig. 4: Distorted tracks found by DeepConcolic testing

the testing result, from the latter, results in less deviation from the original track; such information provides insight to evaluate and improve the tracking system.

Meanwhile, we also observe that by only testing the deep learning component it may not be sufficient to mislead the overall tracking system. As in Figure 5, the tracking finally converges to the original one. This is due to compensation provided by other components in the system, and it answers the research question, Q1, that the tracking system is resilient to some extent towards the deficits of a learning component.

Fig. 5: Adversarial tracks (red) after testing different parts of the original track (green)

Finally, changing more frames does not necessarily result in larger deviation from the original track. This is demonstrated as in Figure 6, where a larger distance between the original and adversarial tracks is found when testing frames x_9, x_{10} , instead of x_9, x_{10}, x_{11} . This observation reflects the uncertainty from other components (other than the CNNs) in the tracking system, as specified in research question Q2.

VI. RELATED WORKS

Verification of DNNs from software testing perspective has been a popular direction [5], [23], [24], [25], [26], [27], [28],

Fig. 6: Adversarial tracks from testing different numbers of consecutive frames starting from x_9

[29], [30], [31], [32], [33], [34], see [1] for a survey with other perspectives such as formal verification and interpretability. Autonomous driving has been the primary application domain for assessments of DNN testing techniques [35], [36], [37], [38]. The analysis in [35], [36] comprises an image generator that produces synthetic pictures for testing neural networks used in classification of cars in autonomous vehicles. In [37], a 3D simulator is proposed to test the dynamics of the pedestrians and the agent vehicles (including simple dynamics for suspension, tyres, etc.), in the virtual environment of the system under test. In [38], a technique is presented to reason about the safety of a closed-loop, learning-enabled control system. Research is also conducted on the verification of cognitive trust between human and autonomous systems [39].

We note that there is limited work focussed on the verification of learning-enabled systems such as the one highlighted in this work.

VII. CONCLUSIONS

In this paper, we show that solely testing the correctness of deep learning components in a vehicle tracking system in isolation is insufficient, since either a deficit discovered in the learning component may be suppressed by the existence of other components, or that there are new uncertainties introduced due to the interaction between learning and non-learning components. Similar results are also observed for the connections between LSTM and CNN layers [40]. These results clearly indicate the necessity of developing a testing strategy that addresses learning components in isolation, as well as in combination with other components within a wider system-level testing framework for learning-enabled systems.

This document is an overview of UK MOD (part) sponsored research and is released for informational purposes only. The contents of this document should not be interpreted as representing the views of the UK MOD, nor should it be assumed that they reflect any current or future UK MOD policy. The information contained in this document cannot supersede any statutory or contractual requirements or liabilities and is offered without prejudice or commitment.

Content includes material subject to © Crown copyright (2018), Dstl. This material is licensed under the terms of the Open Government Licence except where otherwise stated. To view this licence, visit <http://www.nationalarchives.gov.uk/doc/open-government-licence/version/3> or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk.

REFERENCES

- [1] Xiaowei Huang, Daniel Kroening, Wenjie Ruan, Youcheng Sun, Emese Thamo, Min Wu, and Xinpeng Yi. A survey of safety and trustworthiness of deep neural networks. *arXiv preprint arXiv:1812.08342*, 2018.
- [2] Joseph Sifakis. Autonomous systems - an architectural characterization. *CoRR*, abs/1811.10277, 2018.
- [3] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. DeepXplore: Automated whitebox testing of deep learning systems. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 1–18. ACM, 2017.
- [4] Youcheng Sun, Xiaowei Huang, Daniel Kroening, James Sharp, Matthew Hill, and Rob Ashmore. Structural test coverage criteria for deep neural networks. *ACM Transactions on Embedded Computing Systems*, 2019.
- [5] Youcheng Sun, Min Wu, Wenjie Ruan, Xiaowei Huang, Marta Kwiatkowska, and Daniel Kroening. Concolic testing for deep neural networks. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE*, pages 109–119, 2018.
- [6] Kelly Hayhurst, Dan Veerhusen, John Chilenski, and Leanna Rierston. A practical tutorial on modified condition/decision coverage. Technical report, NASA, 2001.
- [7] Yifan Zhou and Simon Maskell. Detecting and tracking small moving objects in wide area motion imagery (WAMI) using convolutional neural networks (CNNs). In *FUSION*, 2019.
- [8] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- [9] Lei Ma, Felix Juefei-Xu, Jiyuan Sun, Chunyang Chen, Ting Su, Fuyuan Zhang, Minhui Xue, Bo Li, Li Li, Yang Liu, Jianjun Zhao, and Yadong Wang. DeepGauge: Comprehensive and multi-granularity testing criteria for gauging the robustness of deep learning systems. In *Automated Software Engineering (ASE)*, pages 120–131. ACM, 2018.
- [10] Youcheng Sun, Xiaowei Huang, Daniel Kroening, James Sharp, Matthew Hill, and Rob Ashmore. Structural test coverage criteria for deep neural networks. In *Proceedings of the 41st International Conference on Software Engineering: Companion Proceedings*, pages 320–321, 2019.
- [11] Xiaowei Huang, Marta Kwiatkowska, Sen Wang, and Min Wu. Safety verification of deep neural networks. In *International Conference on Computer Aided Verification*, pages 3–29. Springer, 2017.
- [12] Matthew Wicker, Xiaowei Huang, and Marta Kwiatkowska. Feature-guided black-box safety testing of deep neural networks. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, *arXiv preprint arXiv:1710.07859*, 2018.
- [13] Wenjie Ruan, Xiaowei Huang, and Marta Kwiatkowska. Reachability analysis of deep neural networks with provable guarantees. *IJCAI*, 2018.
- [14] Wenjie Ruan, Min Wu, Youcheng Sun, Xiaowei Huang, Daniel Kroening, and Marta Kwiatkowska. Global robustness evaluation of deep neural networks with provable guarantees for the hamming distance. *IJCAI*, 2019.
- [15] Min Wu, Matthew Wicker, Wenjie Ruan, Xiaowei Huang, and Marta Kwiatkowska. A game-based approximate verification of deep neural networks with provable guarantees. *Theoretical Computer Science*, 807:298 – 329, 2020. In memory of Maurice Nivat, a founding father of Theoretical Computer Science - Part II.
- [16] AFRL. Wright-patterson air force base (wpafb) dataset. <https://www.sdms.afrl.af.mil/index.php?collection=wpafb2009>, 2009.
- [17] Lars Wilko Sommer, Michael Teutsch, Tobias Schuchert, and Jürgen Beyerer. A survey on moving object detection for wide area motion imagery. In *IEEE Winter Conference on Applications of Computer Vision*, WACV, pages 1–9, 2016.
- [18] Rodney LaLonde, Dong Zhang, and Mubarak Shah. Clusternet: Detecting small objects in large scenes by exploiting spatio-temporal information. In *CVPR*, pages 4003–4012, 2018.
- [19] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, November 2004.
- [20] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In Aleš Leonardis, Horst Bischof, and Axel Pinz, editors, *Computer Vision – ECCV*, pages 404–417, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [21] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: An efficient alternative to sift or surf. In *International Conference on Computer Vision*, pages 2564–2571, Nov 2011.
- [22] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981.
- [23] Youcheng Sun, Xiaowei Huang, Daniel Kroening, James Sharp, Matthew Hill, and Rob Ashmore. DeepConcolic: Testing and debugging deep neural networks. In *Proceedings of the 41st International Conference on Software Engineering*. IEEE, 2019.
- [24] Divya Gopinath, Kaiyuan Wang, Mengshi Zhang, Corina S Pasareanu, and Sarfraz Khurshid. Symbolic execution for deep neural networks. *arXiv preprint arXiv:1807.10439*, 2018.
- [25] Augustus Odena and Ian Goodfellow. TensorFuzz: Debugging neural networks with coverage-guided fuzzing. *arXiv preprint arXiv:1807.10875*, 2018.
- [26] Xiaofei Xie, Lei Ma, Felix Juefei-Xu, Hongxu Chen, Minhui Xue, Bo Li, Yang Liu, Jianjun Zhao, Jianxiang Yin, and Simon See. Coverage-guided fuzzing for deep neural networks. *arXiv preprint arXiv:1809.01266*, 2018.
- [27] Jingyi Wang, Guoliang Dong, Jun Sun, Xinyu Wang, and Peixin Zhang. Adversarial sample detection for deep neural network through model mutation testing. In *Proceedings of the 41st International Conference on Software Engineering*. IEEE, 2019.
- [28] Jingyi Wang, Jun Sun, Peixin Zhang, and Xinyu Wang. Detecting adversarial samples for deep neural networks through mutation testing. *arXiv preprint arXiv:1805.05010*, 2018.
- [29] Lei Ma, Fuyuan Zhang, Jiyuan Sun, Minhui Xue, Bo Li, Felix Juefei-Xu, Chao Xie, Li Li, Yang Liu, Jianjun Zhao, et al. DeepMutation: Mutation testing of deep learning systems. In *Software Reliability Engineering, IEEE 29th International Symposium on*, 2018.
- [30] Dawei Cheng, Chun Cao, Chang Xu, and Xiaoxing Ma. Manifesting bugs in machine learning code: An explorative study with mutation testing. In *International Conference on Software Quality, Reliability and Security (QRS)*, pages 313–324. IEEE, 2018.
- [31] Weijun Shen, Jun Wan, and Zhenyu Chen. MuNN: Mutation analysis of neural networks. In *International Conference on Software Quality, Reliability and Security Companion, QRS-C*. IEEE, 2018.
- [32] Junhua Ding, Xiaojun Kang, and Xin-Hua Hu. Validating a deep learning framework by metamorphic testing. In *Metamorphic Testing (MET), 2nd International Workshop on*, pages 28–34. IEEE, 2017.
- [33] Anurag Dwarakanath, Manish Ahuja, Samarsh Sikand, Raghotham M Rao, RP Bose, Neville Dubash, and Sanjay Podder. Identifying implementation bugs in machine learning based image classifiers using metamorphic testing. In *Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 118–128. ACM, 2018.
- [34] Mengshi Zhang, Yuqun Zhang, Lingming Zhang, Cong Liu, and Sarfraz Khurshid. DeepRoad: GAN-based metamorphic autonomous driving system testing. In *Automated Software Engineering (ASE), 33rd IEEE/ACM International Conference on*, 2018.
- [35] Cumhur Erkan Tuncali, Georgios Fainekos, Hisahiro Ito, and James Kapinski. Simulation-based adversarial test generation for autonomous vehicles with machine learning components. In *Intelligent Vehicles Symposium (IV)*, pages 1555–1562. IEEE, 2018.
- [36] Tommaso Dreossi, Alexandre Donzé, and Sanjit A Seshia. Compositional falsification of cyber-physical systems with machine learning components. In *NASA Formal Methods Symposium*, pages 357–372. Springer, 2017.
- [37] Cumhur Erkan Tuncali, Georgios Fainekos, Hisahiro Ito, and James Kapinski. Simulation-based adversarial test generation for autonomous vehicles with machine learning components. *arXiv preprint arXiv:1804.06760*, 2018.
- [38] Cumhur Erkan Tuncali, Hisahiro Ito, James Kapinski, and Jyotirmoy V Deshmukh. Reasoning about safety of learning-enabled components in autonomous cyber-physical systems. In *55th Design Automation Conference (DAC)*, pages 1–6. IEEE, 2018.
- [39] Xiaowei Huang, Marta Kwiatkowska, and Maciej Olejnik. Reasoning about cognitive trust in stochastic multiagent systems. *ACM Trans. Comput. Logic*, 20(4), July 2019.
- [40] Wei Huang, Youcheng Sun, James Sharp, and Xiaowei Huang. Test metrics for recurrent neural networks. *arXiv preprint arXiv:1911.01952*, 2019.