Monitoring and Diagnosability of Perception Systems

Pasquale Antonante, David I. Spivak, Luca Carlone

Abstract—Perception is a critical component of high-integrity applications of robotics and autonomous systems, such as selfdriving vehicles. In these applications, failure of perception systems may put human life at risk, and a broad adoption of these technologies requires the development of methodologies to guarantee and monitor safe operation. Despite the paramount importance of perception systems, currently there is no formal approach for system-level monitoring. In this work, we propose a mathematical model for runtime monitoring and fault detection and identification in perception systems. Towards this goal, we draw connections with the literature on diagnosability in multiprocessor systems, and generalize it to account for modules with heterogeneous outputs that interact over time. The resulting temporal diagnostic graphs (i) provide a framework to reason over the consistency of perception outputs -across modules and over time- thus enabling fault detection, (ii) allow us to establish formal guarantees on the maximum number of faults that can be uniquely identified in a given perception system, and (iii) enable the design of efficient algorithms for fault identification. We demonstrate our monitoring system, dubbed PerSyS, in realistic simulations using the LGSVL self-driving simulator and the Apollo Auto autonomy software stack, and show that PerSyS is able to detect failures in challenging scenarios (including scenarios that have caused self-driving car accidents in recent years), and is able to correctly identify faults while entailing a minimal computation overhead ($< 5 \,\mathrm{ms}$ on a single-core CPU).

SUPPLEMENTARY MATERIAL

Video: https://youtu.be/GJ9esRS8jZs

I. Introduction

The automotive industry is undergoing a change that could revolutionize mobility. Self-driving cars promise a deep transformation of personal mobility and have the potential to improve safety, efficiency (e.g., commute time, fuel), and induce a paradigm shift in how entire cities are designed [1]. One key factor that drives the adoption of such technology is the capability of ensuring and monitoring safe operation. Consider Uber's fatal self-driving crash [2] in 2018: the report from the National Transportation Safety Board states that "inadequate safety culture" contributed to the fatal collision between the autonomous vehicle and the pedestrian. The lack of safety guarantees, combined with the unavailability of formal monitoring tools, is the root cause of these accidents and has a profound impact on the user's trust. The American Automobile Association's survey 3 shows that 71% of Americans claim to be afraid of riding in a self-driving car. This is a clear sign that the industry needs a sound methodology, embedded in the design process, to guarantee safety and build public trust.

P. Antonante, D. I. Spivak, and L. Carlone are with the Laboratory for Information & Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA. {antonap, dspivak, lcarlone}@mit.edu

This work was partially funded by the NSF CAREER award "Certifiable Perception for Autonomous Cyber-Physical Systems", AFOSR FA9550-20-1-0348 and by MathWorks.





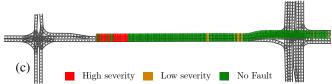


Fig. 1. (a) Tesla Autopilot fails to detect an overturned truck, leading to a crash (Taiwan, 2020) ①. (b) We replicate the same scenario in the LGSVL simulator ②. and observe it still causes accidents in a state-of-the-art open-source autonomy stack (Apollo Auto ③.). (c) We propose PerSyS, a graph-theoretic framework for fault detection and identification, and show it is able to detect failures and prevent crashes in realistic simulations, including the scenario in (b). The colors in the figure show the result of the fault detection.

While safety guarantees have been investigated in the context of control and decision-making [7], [8], the state of the art is still lacking a formal and broadly applicable methodology for monitoring *perception systems*, which constitute a key component of any autonomous vehicle. Perception systems provide functionalities such as localization and obstacle mapping, lane detection, detection and tracking of other vehicles, pedestrians, and traffic signs, among others.

State of Practice. The automotive industry currently uses four classes of methods to claim the safety of an autonomous vehicle (AV) [9], namely: miles driven, simulation, scenario-based testing, and disengagement. Each of these methods has well-known limitations. The *miles driven* approach is based on the statistical argument that if the probability of crashes per mile is lower in autonomous vehicles than for humans, then AVs are safer; however, such an analysis would require an im-

practical amount (i.e., billions) of miles to produce statisticallysignificant results [10], [9]. The same approach can be made more scalable through simulation, but unfortunately creating a life-like simulator is an open problem, for some aspects even more challenging than self-driving itself. Scenario-based testing is based on the idea that if we can enumerate all the possible driving scenarios that could occur, then we can simply expose the AV (via simulation, closed-track testing, or on-road testing) to all these scenarios and, as a result, be confident that the AV will only make sound decisions. However, enumerating all possible corner cases (and perceptual conditions) is a daunting task. Finally, disengagement [11] is defined as the moment when a human safety driver has to intervene in order to prevent a hazardous situation. However, while less frequent disengagements indicate an improvement of the AV behavior, they do not give evidence of the system safety.

An established methodology to ensure safety is to develop a standard that every manufacturer has to comply with. In the automotive industry, the standard ISO 26262 [12] is a risk-based safety standard that applies to electronic systems in production vehicles. A key issue is that ISO 26262 relies on the presence of human drivers (and mostly focuses on electronic systems rather than algorithmic aspects) hence it does not readily apply to fully autonomous vehicles [13]. The recent ISO 21448 [14] is intended to complement ISO 26262, but currently mostly provides high-level considerations to inform stakeholders. The report [15] discusses verification and validation for highly autonomous vehicles, and suggests the use of *monitors* to detect off-nominal performance. Perception has been the subject to increasing attention also outside the automated driving industry [16], [17]. The report [18] considers an automatic aircraft landing system as case study and investigates the challenges in developing trustworthy AI, with focus on machine learning. These reports stress the role of perception for autonomous systems and motivate us to design a rigorous framework for system-level monitoring.

Related Work. Formal methods [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [7] have been recently used as a tool to study safety of autonomous systems. These approaches have been successful for decision systems, such as obstacle avoidance [7], road rule compliance [30] and high-level decision-making [31], where the specifications are usually model-based and have well-defined semantics [8]. However, they are challenging to apply to perception systems, due to the complexity of modeling the physical environment [20], and the trade-off between evidence for certification and tractability of the model [32]. Current approaches [33], [34], [35], [36] consider high-level abstractions of perception [9], [37], [38] or rely on simulation to assert the true state of the world [33], [34], [39].

Relevant to the approach presented is this paper is the class of *runtime verification* methods. Runtime verification is an online approach based on extracting information from

a running system and using it to detect (and possibly react to) observed behaviors satisfying or violating certain properties [40], [41]. Traditionally, the task of evaluating whether a module is working properly is assigned to a monitor, which verifies some input/output properties on the module alone. Balakrishnan *et al.* [37] use Timed Quality Temporal Logic (TQTL) to reason about desiderable spatio-temporal properties of a perception algorithm. Kang *et al.* [42] use model assertions, which similarly place a logical constraint on the output of a module to detect anomalies. We argue that this paradigm does not fully capture the complexity of perception pipelines: while monitors can be used to infer the state of a module, perception pipelines provide an additional opportunity to cross-check the *compatibility* of results across different modules. In this paper we try to address this limitation.

Fault-tolerant architectures [43] have also been proposed to detect and potentially recover from a faulty state, but these efforts mostly focus on implementing watchdogs and monitors for specific modules, rather than providing tools for system-level analysis and monitoring.

Finally, performance guarantees for perception have been investigated for specific problems. In particular, related work on *certifiable algorithms* [44], [45], [46] provides algorithms that are capable of identifying faulty behaviors during execution. These related works mostly focus on specific algorithms, while our goal is to establish monitoring for perception *systems*, including multiple interacting modules and algorithms.

Contribution. In this paper we develop a methodology to detect and identify faulty modules in a perception pipeline at runtime. In particular, we address two questions (adapted from [47]): (i) Can I (as a developer) verify that the perception modules are providing reliable interpretations of the sensor data? (ii) Can I (as regulator) trace the steps that led to an accident caused by an autonomous vehicle? Towards this goal, our first contribution is to draw connections between perception monitoring and the literature on diagnosability in multiprocessor systems, and in particular the PMC model [48]. Our second contribution is to generalize the PMC model to account for modules with heterogeneous outputs, and add a temporal dimension to the problem to account for modules interacting over time. This results in the notion of temporal diagnostic graphs, which (i) provide a framework to reason over the consistency of perception outputs -across modules and over time- thus enabling failure detection, (ii) allow us to establish formal guarantees on the maximum number of faults that can be uniquely identified in a given perception systems. and (iii) enable the design of efficient algorithms for fault identification. Our third contribution is to demonstrate our monitoring system, dubbed PerSyS, in realistic simulations using the LGSVL self-driving simulator and the Apollo Auto autonomy software stack, and show that PerSyS is able to detect failures in challenging scenarios (including night and adverse weather), including conditions that have caused self-driving

¹Moreover, the analysis should cover all representative driving conditions (*e.g.*, driving on a highway is easier than driving in urban environment) and should be repeated at every software update, quickly becoming impractical.

²Our notion of *diagnosability* is related to the level of redundancy within the system and provides a quantitative measure of robustness.

car accidents in recent years (Fig. \blacksquare). Moreover, we show that PerSyS is able to correctly identify faults –whenever the system is sufficiently *diagnosable*– while entailing a minimal computation overhead ($< 5 \, \mathrm{ms}$ on a single-core CPU).

II. MONITORING OF PERCEPTION SYSTEMS

This section presents a mathematical model for monitoring and fault diagnosis in perception systems. Our model *detects* if one of the modules in a perception system is returning an incorrect output, and possibly *identifies* the faulty modules. Section II-A reviews results on fault detection and *diagnosability* in multiprocessors systems, which has been extensively studied since the late 1960s. Sections II-B and II-C extend diagnosability to perception systems and discuss how to model temporal aspects arising in AV applications.

A. Diagnostic Graphs and Diagnosability

Our approach builds on the PMC model [48] for fault diagnosis. In PMC, a set of processors is assembled such that each processor has the capability to communicate with a subset of the other processors, and all of the processors perform the same computation. In such system, a fault occurs whenever the outputs of some processors disagree with each other; the problem is then to identify which processor is faulty. In the PCM model, each processor is assigned a subset of the other processors for the purpose of testing. Using a comparisonbased mechanism, the model aims to characterize the set of faulty processors. Clearly, it is not possible to determine the faulty subset in general, so much of the literature on multiprocessor diagnosis considers two fundamental questions [49]: (i) Given a collection of processors and a set of tests, what is the maximum number of arbitrary processors that can be faulty such that the set of faulty processors can be uniquely identified? (ii) Given a set of test results, does there exist an efficient procedure to identify the faulty processors? The key tool to address these questions is the *diagnostic graph* [48].

Diagnostic Graph. At any given time, each processor is assumed to be in one of two states: *faulty* or *fault-free*. Diagnosis is based on the ability of processors to test—*i.e.*, to check the consistency of its output against— other processors. Formally we assume that each processor implements one or more *consistency functions*; these are Boolean functions that return *pass* (0) or *fail* (1), depending on whether the output of two processors is in agreement.

To perform the diagnosis, we follow [48] and model the problem as a directed graph D=(U,E), where U is the set of processors, while the edges E represent the test assignments. In particular, for an edge $(i,j) \in E$, we say that node i is testing node j. The outcome of this test is the result of the consistency functions of i, in other words 1 (resp. 0) if i evaluates j as faulty (resp. fault-free). Fault-free processors are assumed to provide correct test results, whereas no assumption is made about tests executed by faulty processors: they may produce correct or incorrect test outcomes. We call D the diagnostic graph. The collection of all test results for a test assignment E is called a syndrome. Formally, a syndrome is a

function $\sigma: E \to \{0,1\}$. The syndrome is then processed to make some determination about the faultiness status of every processor in the system. The notion of κ -diagnosability formalizes when it is possible to use the syndrome to diagnose and identify all faults in a system.

Definition 1 (κ -diagnosability [48]). A diagnostic graph D = (U, E) with |U| processors is κ -diagnosable (with $\kappa < |U|$) if, given any syndrome, all faulty processors can be identified, provided that the number of faults does not exceed κ .

The problem of determining the maximum value of κ for which a given system is κ -diagnosable is called the diagnosability problem. We denote the maximum value of κ -diagnosability of a graph D by $\kappa(D)$.

Characterization of κ -diagnosability. Consider a directed graph D=(U,E). For $i\in U$, we denote with $\delta_{\mathrm{in}}(i)$ the number of edges directed toward i (in-degree). We denote by $\delta_{\mathrm{in}}(D)=\min_{i\in U}d_{\mathrm{in}}(i)$ the minimum in-degree of the graph. We denote with $\Gamma(i)=\{j\in U\mid (i,j)\in E\}$ the *testable* set of i (outgoing neighbors of i). Finally, for $X\subset U$, we define $\Gamma(X)=\{\bigcup_{i\in X}\Gamma(i)\setminus X\}$. We can now state the following theorem that shows how to check if a graph is κ -diagnosable.

Theorem 2 (Characterization of κ -diagnosability [50]). Let D = (U, E) be a diagnostic graph with |U| processors. Then D is κ -diagnosable if and only if

- (i) $\kappa \leq \frac{|U|-1}{2}$;
- (ii) $\kappa \leq \bar{\delta_{in}(D)}$; and
- (iii) for each integer p with $0 \le p < \kappa$, and each $X \subset U$ with $|X| = |U| 2\kappa + p$ we have $|\Gamma(X)| > p$.

A naive implementation of the checks in Theorem 2 would lead to an algorithm of time complexity $O(|U|^{t+2})$ [51]. Bhat [51] proposes an improved algorithm to fine maximum κ of an arbitrary graph in $O(|U|^{u+2}\log_2(u))$ where $u=\min\{\delta_{\rm in}(D),\lfloor(|U|-1)/2\rfloor\}$. Since this approach can be impractical for large graphs, they also propose a polynomial-time algorithm to find a lower-bound for κ in $O(|U|^{3/2}+|E|)$.

Fault Identification. Once we have a syndrome on a κ diagnosable graph, we would like to actually identify the set of faulty processors, provided that the number of faults does not exceed κ . Dahbura et al. [49] show that the problem of identifying the set of faulty processors is related to the problem of finding the minimum vertex cover set of an undirected graph. In general, the problem of finding a minimum vertex cover is in the class of NP-complete problems, meaning that there is no known deterministic algorithm that is guaranteed to solve the problem in polynomial time, but the validity of any solution can be tested in polynomial time. However, the work [49] exploits some special properties of κ -diagnosable graphs to propose an algorithm with time complexity $O(|U|^{2.5})$ for fault identification. Sullivan [52] proposes an algorithm with time complexity $O(\kappa^3 + |E|)$ for fault identification and proves this is the tightest bound if κ is $o(|U|^{5/6})$. These results reassure us that the framework of diagnostic graphs provides a practical basis to build our real-time perception system monitor.

B. Temporal Diagnostic Graphs

The PMC model focuses on instantaneous fault diagnosis. In other words, given a diagnostic graph and a syndrome at time t, PMC-based algorithms detect and possibly identify faults only considering tests occurring at that time instant. However, perception modules evolve over time, are asynchronous, and output data at different rates. In the following, we extend the notion of diagnostic graphs to account for the temporal dimension of perception. Considering the temporal dimension will allow us to check consistency over time (we describe the temporal consistency tests in Section Π -C).

Intuitively, if we consider an arbitrary time interval [a,b], over the interval we have the opportunity to collect multiple outputs for each module in a perception system. Furthermore, the outputs of these modules cannot be arbitrary, and must have some temporal consistency. For instance, because the vehicle is moving at a finite velocity, measurements from the GPS at consecutive times should be within a maximum distance. Similarly, since external objects detected by the perception system are stationary or move at bounded speed, consecutive detections must not be too spaced apart. Therefore, we can build a diagnostic graph over the interval [a,b] and test consistency of the perception results across modules and over time, and detect failures when these tests fail.

More formally, we define an [a,b]-temporal diagnostic graph $D_{[a,b]}$ as a directed graph where each node represents the execution of a module in the system within the time interval [a,b]. Each module (e.g.), sensor fusion, object detection, LIDAR sensor) can be executed multiple times in [a,b], hence inducing multiple nodes in $D_{[a,b]}$. Moreover, the graph has edges whenever an arbitrary pair of modules can test (instantaneously or over time) their inputs or outputs. Intuitively, one can think at a temporal diagnostic graph as a larger graph where we "stack" multiple diagnostic graphs for each time $t \in [a,b]$, and add temporal tests.

Besides being more expressive, temporal diagnostic graph can improve diagnosability. We formalize this aspect by first observing that adding edges (*e.g.*, temporal tests) in the graph can only improve diagnosability (all proofs in Appendix).

Corollary 3. If $D \subseteq D'$ are two graphs with the same set of vertices, then $\kappa(D) \leq \kappa(D')$.

Then we observe that connecting multiple κ -diagnosable graphs over time cannot decrease the diagnosability.

Proposition 4. Let D = (U, E) be a diagnostic graph. Given subsets $U_i \subseteq U$ that cover U (in the sense that $U = \bigcup_{i \in I} U_i$), and for each $i \in I$, let $D_i \subseteq D$ be the largest subgraph with vertices U_i . Then if each D_i is κ -diagnosable, so is D.

In the experimental section, we will observe that if the temporal tests are well-designed, the diagnosability strictly increases when considering larger time intervals.

C. Consistency Functions in Perception Systems

Each edges in our temporal diagnostic graph implements a consistency function to test consistency between the processing at different modules in the perception system. This section defines four classes of consistency tests for a perception system with arbitrary modules interacting over time.

- 1) Input Admissibility: In this type of consistency test, a module i monitors its inputs \mathcal{I} to judge whether they are admissible (e.g., the IMU measurements fall within an admissible range). This in turns induces a consistency test between module i and another module j producing \mathcal{I} . Admissible ranges are known a priori in many perception modules.
- 2) Input or Output Consistency: In this type of test, a module monitors the consistency of inputs coming from (or outputs produced by) multiple other modules. For instance, consider two modules i and j performing object detection using two different sensors, e.g., camera and LIDAR. Both modules output a list of obstacles and, by comparison, i can verify that j is producing a consistent detection. This class of tests is common in perception, due to sensor redundancy.
- 3) Input/Output Consistency: In this class we include all tests implemented by sensor fusion modules, which test their fused estimate against the input data. Input/Output Consistency differs from Input Consistency in that it needs to compute the output variable in order to test the validity of the inputs. As an example, consider the case in which Visual Odometry (module j) incorrectly estimates the motion of the vehicle. After fusing multiple incoming odometry measurements, the sensor fusion module i can identify that visual odometry produced an incorrect measurement [53], [54], [55].
- 4) Temporal Consistency: In this class we include all the tests between a module i and a module j involving data produced at different timestamps. The simplest form of temporal consistency involves an edge connecting two nodes corresponding to the outputs of the same module at consecutive time steps, e.g., i and j being two consecutive IMU measurements or pose estimates. The tests can be extended to involve different modules and non-consecutive timestamps.

D. Discussion

The diagnostic graph is a directed graph. An edge (i,j) indicates that module i tests j, reporting if the output of j is correct from the standpoint if i. The direction of the edge is useful whenever the test is non-symmetric, for example when a module (e.g., a watchdog) is testing the output of another module (e.g., state estimation) but not vice-versa. When the test is symmetric (i.e., two modules test if their outputs are compatible with each other), both edge directions are theoretically valid (in our tests, we assign an arbitrary direction to the edge). The binary outcome of a consistency test asserts consistency across the modules' outputs.

The perception system experiences a fault whenever two modules produce an inconsistent representation of the environment. The faults can be intermittent, that is, occur at one time and disappear at the next. Moreover, a consistency test can implement an arbitrary logic to detect an inconsistency as long as it returns a boolean value indicating whether it detected an inconsistency (fault) or not. For instance, consider an object detection system including camera and LIDAR, where each

detected object has a confidence level associated with it. The consistency test is tasked with determining if two sets of obstacles detected by the two sensors are in agreement. The LIDAR might detect an object on the left of the car while the front-facing camera fails to do so. This particular case may not be signaled as inconsistency if the object seen by the LIDAR is outside the camera's field of view. Similarly, if the camera detects an obstacle with low confidence, while the LIDAR does not detect it, the consistency test might be setup not to signal the inconsistency as the confidence is low.

The use of consistency functions generalizes the concept of *voting* in estimation (*e.g.*, [56]), enabling heterogeneous tests and arbitrarily complex logics to detect inconsistencies. For example, while a consistency function verifying that a bounding box associated with a pedestrian occupies a meaningful region of the map, *i.e.*, sidewalks instead of the middle of a highway, can be easily incorporated into our system, it may be hard to include in standard estimation frameworks.

Finally, we note that our framework can also deal with the case where a module produces a correct output despite using an incorrect input. For instance, several robust estimation schemes (e.g., [54]) can produce accurate results despite being fed with incorrect measurements.

III. EXPERIMENTAL EVALUATION

This section shows that temporal diagnostic graphs are effective to detect and identify failures in perception systems in real-time. We validate our theory in both numerical tests (Section III-A) and self-driving car simulations (Section III-B).

Implementation Details. We build temporal diagnostic graphs in Python and perform fault identification using a constraint programming implementation of [52], [49] using Google OR-Tools [57]. The code is executed on a Linux machine with an Intel i-97920X processor (4.3 GHz) and two NVIDIA GeForce GTX 1080 Ti graphics cards. We call the resulting system PerSyS (PERception SYstem Supervisor). The specific choice of consistency tests is discussed below.

A. Fault Identification in Diagnostic Graphs

We start by evaluating PerSyS's capability of identifying faults. We perform a Monte Carlo analysis over randomly generated temporal diagnostic graphs. To test the accuracy of the fault identification, measured as the percentage of correctly identified faults, we generate random 5-diagnosable graphs with 15 nodes and collect the fault identification results for increasing number of faults. For each tested number of faults, we average the results over 100 random graph. Fig. 2a shows that PerSyS is able to correctly identify all the faults whenever their number is below the theoretical limit of 5 (remember that we simulate 5-diagnosable graphs); in the presence of more than 5 faults, PerSyS is still able to detect a failure, but its capability of *identifying* which modules are faulty decreases for increasing amounts of faults. Fig. 2b considers a 5-diagnosable graphs with 5 randomly generated faults and evaluates the CPU time for increasing number of nodes (from 11 to 25). The time required to identify the faults grows

linearly with the number of nodes and does not exceed $5\,\mathrm{ms}$ for graphs with up to 25 nodes.

B. Fault Detection in AV Perception Systems

We use PerSyS to monitor a state-of-the-art perception system, implemented in the open-source Baidu's Apollo Auto 6 autonomous driving platform. We test the system on the Lincoln MKZ car, simulated using LG's LGSVL 5 driving simulator. In our experiments we run on the same Linux machine the simulator, the AV software stack, and PerSyS in real-time. We test PerSyS in object detection and vehicle localization problems and show it is able to prevent accidents in the majority of the tests, even in adverse weather conditions. Following a functional system architectures 58, which separates the perception system from the decision system, when PerSyS identifies a fault in the perception system, it notifies the decision system that is in charge of reacting to it.

1) 3D Object Detection: We use PerSyS to detect failures in the subsystem in charge of detecting and localizing objects around the AV (e.g., other vehicles, pedestrians).

Implementation details. The Lincoln MKZ car is equipped with a GPS, IMU, RADAR, LIDAR (32 channels), a frontfacing monocular camera used for obstacle detection and a telephoto camera (pointing slightly upward) used only for traffic light detection. The Apollo Auto perception system has been configured to use RADAR, LIDAR, and the front-facing camera to detect objects. Apollo Auto's object detection system outputs a list of objects for each sensor and then it fuses each individual detection into a single estimate using a configurable sensor fusion algorithm [59]. In the case of RADAR and camera, each object is parametrized with its 3D bounding box. In the case of the LIDAR each object is represented by its 3D convex hull. We denote with \mathcal{O}_{C} , \mathcal{O}_{L} , and \mathcal{O}_{R} the obstacles detected by the camera, LIDAR, and RADAR, respectively, and with \mathcal{O}_{F} the obstacle as determined by the sensor fusion module (represented by convex hulls). From Theorem 2 we know that a diagnostic graph with 4 nodes can be up to 1-diagnosable; therefore, we implemented several consistency tests, which are visualized in Fig. 3a, to achieve a 1-diagnosable graph. A formal description of each test is given in Appendix, while here we provide a high-level overview. For a consistency check $A \rightarrow B$ (node A testing node B), we return 1 (fail) if there exists an object in A within the field-of-view of B that is not matched by any object in B or vice versa, or 0 (pass) otherwise. We also add temporal tests: for each time t, we test outputs at consecutive timestamps *i.e.*, $A_{t-1} \rightarrow A_t$ (consecutive object detections). When possible, we also cross-check executions of different modules at nearby timestamps, i.e., $A_{t-1} \rightarrow B_t$ (consecutive detections across different sensors). The result is the 2-diagnosable graph shown in Fig. 3a. Note that if we neglect the time dimension, the graph becomes 1-diagnosable, remarking the importance of our extension in Section II-B. Due to the asynchronous nature of the perception system, a new diagnostic graph is created as soon as all the nodes in the graph are available, refreshing old values when new ones are available, and resulting in a

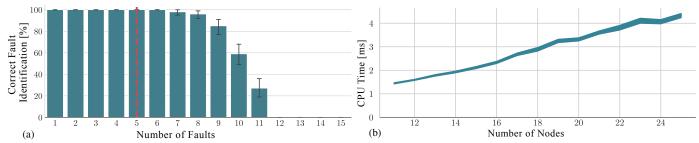


Fig. 2. **Diagnosability and fault identification in temporal diagnostic graphs.** (a) Percentage of correctly identified faults for randomly generated 5-diagnosable graphs with 15 nodes; (b) CPU time with 5 faults and increasing number of nodes. The plots also show the $1-\sigma$ standard deviation.

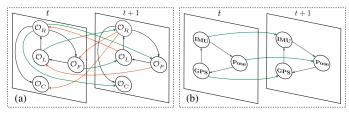


Fig. 3. **Temporal diagnostic graphs.** (a) Object detection system, and (b) Vehicle localization system.

new diagnostic graph every $0.3\text{-}0.5\,\mathrm{ms}$. In our implementation, every consistency test also returns a severity level. For object detection, we set this value to *low* if there is an inconsistency for an object on the road (but not in the current lane), and *high* if there is an inconsistency on the same lane occupied by the AV (Fig. Π_c).

Scenario. In the LGSVL simulator we developed a scenario where an overturned truck occupies a lane in the Borregas Avenue map (Fig. []b). This scenario is designed based on the recent Tesla accident in Taiwan [4], shown in Fig. []a. We tested the system in four different variations of the scenario: (i) Simple scenario takes place at noon with clear weather and no traffic, (ii) Rain takes place at noon but with moderate rain, fog and no traffic, (iii) Night takes places at 9pm with clear weather and no traffic, (iv) Traffic takes place at noon with clear weather and incoming traffic in neighboring lanes. We average results over 10 random instances of each scenario, 5 with the truck occupying a lane and 5 without. Traffic and truck position are randomized, but the truck (when present) always blocks the lane the AV is driving on.

Results. In our tests, Apollo Auto's perception system failed to detect the truck in 70% of the cases, leading to 14 collisions. Fig. To shows the top view of a sample trajectory and also visualizes PerSyS's fault detections. The red sections represent points where a high severity fault was detected, corresponding to the case were the truck was blocking the road but the camera failed to detect it, inducing failures in the sensor fusion module and resulting in a collision. By detecting the failure, PerSyS would be able to trigger an emergency stop and prevent the accident.

Table [I] shows statistics describing PerSyS's effectiveness in detecting failures. The "correct critical" column indicates the percentage of case where PerSyS detected high severity faults, and using it would have prevented an imminent collision (within 20 seconds). We observe that PerSyS would have prevented nearly all accidents in our tests. The "false

non-critical" column reports the percentage of cases where PerSyS triggered low severity faults that did not resolve in a collisions; these are typically caused by the Apollo Auto camera-based object detection, which mis-detected obstacles on the sidewalks or other lanes. These are frequent in modern perception systems, and are classified as non critical by PerSyS. The "Traffic" scenario exhibits the highest number of false non-critical faults: they are the result of the camera failures in the proximity of an incoming vehicle. The "false critical" column describes the percentage of tests where PerSyS would have stopped the car even if no accident was imminent; this percentage is low and can be further improved upon (Section IV). The "Rain" scenario has the largest number of false critical alarms, because in few occasions Apollo Auto's perception system detected regions of the wet road floor as obstacles, hence triggering a fault detection and incorrectly stopping the vehicle. The experiment shows that even if PerSyS may be conservative, it reliably reports high severity faults. In Section IV we discuss possible extensions to this work that can better cope with false alarms.

2) Vehicle localization: In this section, we use PerSyS to detect failures in Apollo Auto's AV localization system.

Implementation details. The car moves at a maximum speed of 45 mph and a default cruise speed of 25 mph. The Apollo Auto vehicle localization system is configured to use the GNSS RTK module. We denote with GPS and IMU the output of the two sensors while we denote as "Pose Estimate" the estimate produced by Apollo Auto's sensor fusion algorithm. We implemented 6 consistency tests, 3 regular consistency tests and 3 temporal consistency tests. The pose estimate and the GPS measurement are compared directly and the test returns 1 (fail) if the difference in the positions is above 1 m threshold. The comparison with the IMU verifies that the change in the pose estimate is compatible with the measured acceleration and angular velocities (Input/Output Consistency) and within the physical limits of the car (Input Admissibility).

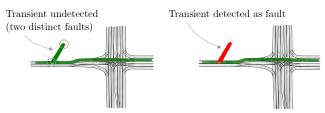
Scenario	False non critical (%)	False critical (%)	Correct critical (%)	Avg. time to impact from failure (s)
Simple	43	0	100	4.17
Night	4.44	0.58	99.42	3.65
Rain	49.55	8.47	91.53	4.04
Traffic	67.67	5.04	94.96	4.11

TABLE I
OBJECT DETECTION RESULTS AGGREGATED BY SCENARIO.

As temporal tests, we compare the output of each module at consecutive time instants, *i.e.*, $A_{t-1} \rightarrow A_t$. The resulting temporal diagnostic graph is 1-diagnosable and visualized in Fig. 3b. A new temporal diagnostic graph is built every $0.1\,\mathrm{ms}$ and we do not distinguish the faults by severity.

Scenario. To test the localization system we used the same scenarios of the previous section, but randomly spoiled GPS measurements with incorrect data. The duration of the fault was randomized to last between 0.1 and 0.4 seconds.

Results. In all the tests, Apollo Auto was able to stop the car without accidents in response to the incorrect GPS readings. We tested PerSyS with a single diagnostic graph and with a temporal diagnostic graph made by 2 consecutive diagnostic graphs (Fig. 3b). PerSyS was also able to detect the failures in all tests, and it correctly identified the GPS and the Pose Estimate as faulty 46.43% of the times in the case of the temporal diagnostic graphs (Fig. 3b). In the case of a single diagnostic graph, PerSyS failed to identify the faulty modules 100% of the times, remarking the importance of using temporal information for monitoring. Incorrect fault identification is due to the fact that when both the GPS and the sensor fusion fail, we end up with two faults in a 1-diagnosable temporal graph, which cannot be uniquely identified. By equipping the sensor fusion with an outlier detector, it would be possible to produce a correct pose estimation using only the IMU; in such a case, the system would have experienced one fault in a 1-diagnosable graph, making fault identification possible. Fig. 4 overlays the output of PerSyS on the vehicle trajectory, comparing traditional and temporal diagnostic graphs. This example strongly indicates the need of using additional sources of information, like visual-inertial odometry, to increase redundacy in the system and hence increase diagnosability.



Using a single diagnostic graph

Using a temporal diagnostic graph

Fig. 4. **Vehicle localization monitoring.** Comparison between diagnostic graph and temporal diagnostic graph for vehicle localization (red: PerSyS detected a fault, green: PerSyS did not detect any faults).

IV. CONCLUSION

We presented PerSyS, a novel framework for fault diagnosis in perception systems. PerSyS reasons over the consistency of the perception results across modules and over time and is shown to be able to detect failures (and prevent accidents) in state-of-the-art perception systems. PerSyS builds on the distant literature of multiprocessor diagnosability, which we extend and tailor to modern perception systems. In particular, we borrow the notion of *diagnosability*, which can be understood as a measure of robustness and quantifies the maximum number of faults that can be correctly identified in a system. Moreover, we extend existing work and introduce *temporal*

diagnostic graphs, which describe perception modules interacting at different rates and over time, while still enabling the use of existing fault identification algorithms.

We believe that *temporal diagnostic graphs* can complement the literature and enhance the current practice, contributing to the goal of achieving safe and trustworthy autonomous vehicles. For instance, a system designer can use the tools proposed in this paper to assess the diagnosability before deploying the vehicle on public roads, or design the system in such a way that its diagnosability is maximized. At runtime, PerSyS allows the vehicle to have a greater awareness of its operational envelope and enables real-time perception monitoring; PerSyS's consistency tests formalize and generalize watchdogs used in standard practice. For regulators, a certain level of diagnosability can be used as a requirement for certification, and, in the unfortunate case of an accident, the proposed approach may increase accountability by providing formal evidence about the root cause of the failures.

In future work, we plan to extend PerSyS to be a probabilistic framework that accounts for, and reasons over, the probability of failure. Such a framework would potentially allow coping with missing (due to asynchronous updates) and unreliable tests, and better model fault propagation.

REFERENCES

- G. Silberg, R. Wallace, G. Matuszak, J. Plessers, C. Brower, and D. Subramanian, "Self-driving cars: The next revolution," White paper, KPMG LLP & Center of Automotive Research, vol. 9, no. 2, pp. 132– 146, 2012.
- [2] NTSB, "Collision between vehicle controlled by developmental automated driving system and pedestrian, tempe, arizona," 2018. [Online]. Available: https://www.ntsb.gov/investigations/AccidentReports/ Reports/HAR1903.pdf
- [3] American Automobile Association, "Three in four americans remain afraid of fully self-driving vehicles," https://newsroom.aaa.com/2019/03/americans-fear-self-driving-cars-survey/, 2019.
- [4] Taiwan English News, "Tesla on autopilot crashes into overturned truck," 2020. [Online]. Available: https://taiwanenglishnews.com/ tesla-on-autopilot-crashes-into-overturned-truck/
- [5] LG, "LGSVL Simulator." [Online]. Available: https://www.lgsvlsimulator.com
- [6] Baidu, "Apollo Auto." [Online]. Available: https://apollo.auto/
- [7] S. Mitsch, K. Ghorbal, D. Vogelbacher, and A. Platzer, "Formal verification of obstacle avoidance and navigation of ground robots," *The International Journal of Robotics Research*, vol. 36, no. 12, pp. 1312–1340, 2017.
- [8] M. Foughali, B. Berthomieu, S. Dal Zilio, P.-E. Hladik, F. Ingrand, and A. Mallet, "Formal verification of complex robotic systems on resourceconstrained platforms," in 2018 IEEE/ACM 6th International FME Workshop on Formal Methods in Software Engineering (FormaliSE), 2018, pp. 2–9.
- [9] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars," ArXiv, vol. abs/1708.06374, 2017.
- [10] N. Kalra and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?" *Transportation Research Part A: Policy and Practice*, vol. 94, pp. 182– 193, 2016.
- [11] I. C. Vehicles, "Google reports self-driving car disengagements," 2015.
 [Online]. Available: https://site.ieee.org/connected-vehicles/2015/12/15/
- 12] ISO Standard, "Road vehicles functional safety," 2011, iSO 26262-1:2011.
- [13] P. Koopman and M. Wagner, "Challenges in autonomous vehicle testing and validation," SAE Int. J. Trans. Safety, vol. 4, no. 1, 2016.
- [14] ISO Standard, "Road vehicles safety of the intended functionality," 2019, iSO/PAS 21448:2019(en).

- [15] Aptiv, Audi, B. Apollo, BMW, Continental, Daimler, F. Group, Here, Infineon, Intel, and Volkswagen, "Safety First for Automated Driving," 2019. [Online]. Available: https://www.daimler.com/innovation/case/autonomous/safety-first-for-automated-driving-2.html
- [16] Uber, "Fast-Forwarding to a Future of On-Demand Urban Air Transportation," 2016. [Online]. Available: https://dlnyezhlys8wfo.cloudfront.net/static/PDFs/Elevate%2BWhitepaper.pdf
- [17] EHang, "Whitepaper on urban air mobility systems," 2020. [Online]. Available: https://www.ehang.com/app/en/EHang%20White%20Paper% 200n%20Urban%20Air%20Mobility%20Systems.pdf
- [18] EASA and Daedalean, Concepts of Design Assurance for Neural Networks, 2020, [pdf].
- [19] F. Ingrand, "Recent trends in formal validation and verification of autonomous robots software," in 2019 Third IEEE International Conference on Robotic Computing (IRC), 2019, pp. 321–328.
- [20] S. Seshia and D. Sadigh, "Towards verified artificial intelligence," ArXiv, vol. abs/1606.08514, 2016.
- [21] X. Ding, B. Englot, A. Pinto, A. Speranzon, and A. Surana, "Hierarchical multi-objective planning: From mission specifications to contingency management," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014, pp. 3735–3742.
- [22] A. Desai, T. Dreossi, and S. Seshia, "Combining model checking and runtime verification for safe robotics," in *International Conference on Runtime Verification*. Springer, 2017, pp. 172–189.
- [23] B. Hoxha and G. Fainekos, "Planning in dynamic environments through temporal logic monitoring," in AAAI Workshop: Planning for Hybrid Systems, vol. 16, 2016, p. 12.
- [24] C.-I. Vasile, J. Tumova, S. Karaman, C. Belta, and D. Rus, "Minimum-violation scLTL motion planning for mobility-on-demand," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2017, pp. 1481–1488.
- [25] S. Dathathri and R. Murray, "Decomposing GR(1) games with singleton liveness guarantees for efficient synthesis," arXiv, vol. abs/1709.07094, 2017
- [26] S. Ghosh, D. Sadigh, P. Nuzzo, V. Raman, A. Donzé, A. Sangiovanni-Vincentelli, S. Sastry, and S. Seshia, "Diagnosis and repair for synthesis from signal temporal logic specifications," in *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '16. ACM, 2016, pp. 31–40.
- [27] W. Li, L. Dworkin, and S. A. Seshia, "Mining assumptions for synthesis," in Ninth ACM/IEEE International Conference on Formal Methods and Models for Codesign (MEMPCODE2011), 2011, pp. 43–50.
- [28] W. Li, D. Sadigh, S. Sastry, and S. Seshia, "Synthesis for human-in-the-loop control systems," in *Intl. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2014.
- [29] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Trans. on Automatic Control*, vol. 53, no. 1, pp. 287–297, 2008.
- [30] N. Roohi, R. Kaur, J. Weimer, O. Sokolsky, and I. Lee, "Self-driving vehicle verification towards a benchmark," arXiv preprint arXiv:1806.08810, 2018.
- [31] R. C. Cardoso, M. Farrell, M. Luckcuck, A. Ferrando, and M. Fisher, "Heterogeneous verification of an autonomous curiosity rover," pp. 353–360, 2020.
- [32] M. Luckcuck, M. Farrell, L. A. Dennis, C. Dixon, and M. Fisher, "Formal specification and verification of autonomous robotic systems: A survey," ACM Computing Surveys (CSUR), vol. 52, no. 5, pp. 1–41, 2019.
- [33] T. Dreossi, D. Fremont, S. Ghosh, E. Kim, H. Ravanbakhsh, M. Vazquez-Chanlatte, and S. Seshia, "VERIFAI: A toolkit for the design and analysis of artificial intelligence-based systems," ArXiv:1902.04245, 2019.
- [34] D. J. Fremont, T. Dreossi, S. Ghosh, X. Yue, A. L. Sangiovanni-Vincentelli, and S. A. Seshia, "Scenic: a language for scenario specification and scene generation," in *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2019, pp. 63–78.
- [35] S. Jha, V. Raman, D. Sadigh, and S. Seshia, "Safe autonomy under perception uncertainty using chance-constrained temporal logic," *Journal* of Automated Reasoning, vol. 60, pp. 43–62, 2017.
- [36] K. Leahy, E. Cristofalo, C. Vasile, A. Jones, E. Montijano, M. Schwager, and C. Belta, "Control in belief space with temporal logic specifications using vision-based localization," *Intl. J. of Robotics Research*, vol. 38, 04 2019.

- [37] A. Balakrishnan, A. G. Puranic, X. Qin, A. Dokhanchi, J. V. Deshmukh, H. Ben Amor, and G. Fainekos, "Specifying and evaluating quality metrics for vision-based perception systems," in *Design, Automation Test* in Europe Conference Exhibition (DATE), 2019, pp. 1433–1438.
- [38] A. Dokhanchi, H. B. Amor, J. Deshmukh, and G. Fainekos, "Evaluating perception systems for autonomous vehicles using quality temporal logic," in *Intl. Conf. on Runtime Verification (RV)*, 2018.
- [39] T. Dreossi, S. Ghosh, A. Sangiovanni-Vincentelli, and S. Seshia, "Systematic testing of convolutional neural networks for autonomous driving," ArXiv, vol. abs/1708.03309, 2017.
- [40] E. Bartocci and Y. Falcone, Lectures on Runtime Verification: Introductory and Advanced Topics. Springer, 2018, vol. 10457.
- [41] A. Francalanza, J. A. Pérez, and C. Sánchez, "Runtime verification for decentralised and distributed systems," in *Lectures on Runtime* Verification. Springer, 2018, pp. 176–210.
- [42] D. Kang, D. Raghavan, P. Bailis, and M. Zaharia, "Model assertions for debugging machine learning," in NIPS, 2018.
- [43] A. Santamaria-Navarro, R. Thakker, D. D. Fan, B. Morrell, and A. akbar Agha-mohammadi, "Towards resilient autonomous navigation of drones," 2020.
- [44] H. Yang and L. Carlone, "A polynomial-time solution for robust registration with extreme outlier rates," in *Robotics: Science and Systems* (RSS), 2019, [opdf], [oideo], [media], [media], [media].
- [45] ——, "In perfect shape: Certifiably optimal 3D shape reconstruction from 2D landmarks," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020, arxiv version: 1911.11924, [pdf].
- [46] J. Briales, L. Kneip, and J. Gonzalez-Jimenez, "A certifiably globally optimal solution to the non-minimal relative pose problem," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [47] M. Brundage, S. Avin, J. Wang, H. Belfield, G. Krueger, G. Hadfield, H. Khlaaf, J. Yang, H. Toner, R. Fong et al., "Toward trustworthy ai development: Mechanisms for supporting verifiable claims," arXiv preprint arXiv:2004.07213, 2020.
- [48] F. P. Preparata, G. Metze, and R. T. Chien, "On the connection assignment problem of diagnosable systems," *IEEE Transactions on Electronic Computers*, no. 6, pp. 848–854, 1967.
- [49] A. T. Dahbura and G. M. Masson, "An $O(n^{2.5})$ fault identification algorithm for diagnosable systems," *IEEE Transactions on Computers*, no. 6, pp. 486–492, 1984.
- [50] S. L. Hakimi and A. T. Amin, "Characterization of connection assignment of diagnosable systems," *IEEE Transactions on Computers*, vol. 100, no. 1, pp. 86–88, 1974.
- [51] K. Bhat, "Algorithms for finding diagnosability level and t-diagnosis in a network of processors," in *Proceedings of the ACM'82 conference*, 1982, pp. 164–168.
- [52] G. F. Sullivan, "An $O(t^3 + |E|)$ fault identification algorithm for diagnosable systems," *IEEE Transactions on Computers*, vol. 37, no. 4, pp. 388–397, 1988.
- [53] P. Lajoie, S. Hu, G. Beltrame, and L. Carlone, "Modeling perceptual aliasing in SLAM via discrete-continuous graphical models," *IEEE Robotics and Automation Letters (RA-L)*, 2019, extended ArXiv version: (pdf), Supplemental Material: (pdf).
- [54] P. Antonante, V. Tzoumas, H. Yang, and L. Carlone, "Outlier-robust estimation: Hardness, minimally tuned algorithms, and applications," *IEEE Trans. Robotics*, 2021, arXiv preprint arXiv: 2007.15109, [cpdf].
- [55] H. Yang and L. Carlone, "One ring to rule them all: Certifiably robust geometric perception with outliers," in Conference on Neural Information Processing Systems (NeurIPS), H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 18 846–18 859, [pdf], [Online]. Available: https://proceedings.neurips.cc/paper/2020/file/da6ea77475918a3d83c7e49223d453cc-Paper.pdf
- [56] H. Yang, J. Shi, and L. Carlone, "TEASER: Fast and Certifiable Point Cloud Registration," *IEEE Trans. Robotics*, vol. 37, no. 2, pp. 314–333, 2020, extended arXiv version 2001.07715 [cpdf].
- [57] Google, "Google OR-Tools." [Online]. Available: https://developers.
- [58] Ö. Ş. Taş, F. Kuhnt, J. M. Zöllner, and C. Stiller, "Functional system architectures towards fully automated driving," in 2016 IEEE Intelligent vehicles symposium (IV). IEEE, 2016, pp. 304–309.
- [59] Baidu, "Apollo Auto." [Online]. Available: https://github.com/ ApolloAuto/apollo

APPENDIX A CONSISTENCY TESTS

We list all the consistency tests used in the experimental evaluation with Apollo Auto.

A. Notation

Definition 1 (Approximate set membership). With $x \in S$ we denote the approximate set membership. We assume there exist a distance function $\mathbf{dist}_S(\cdot,\cdot)$ and a given threshold τ_S such that, if $x \in S$ we have

$$\exists y \in S \text{ s.t. } \mathbf{dist}_S(x,y) \leq \tau_S$$

When we use the approximate set membership to determine if two obstacles are approximately the same, the distance function depends on the data structure representing them while the threshold τ_S depends on their size.

Definition 2 (Objects in Field of View). Given a sensor A we denote with FoV(A) the set of obstacles that fall in the field of view of the sensor.

Definition 3 (Map and Current Lane). We denote with Map the set of points that belong to the map. Similarly, we denote with CurrentLane the set of points that belong to the lane currently occupied by the AV.

When determining if an obstacle belongs to the Field of View of a sensor and/or a lane, we use the approximate set membership operator. Moreover, we denote:

- ullet \mathcal{O}_{R} the set of objects detected by the RADAR
- \bullet \mathcal{O}_{L} the set of objects detected by the LIDAR
- ullet \mathcal{O}_{C} the set of objects detected by the camera
- \mathcal{O}_{F} the set of objects returned by the sensor fusion alg.
- τ_p the maximum position error allowed (30 cm)
- τ_v the maximum velocity error allowed $(0.5\,\mathrm{m/s})$
- \hat{v} the maximum vehicle speed (90 km/h)
- \hat{a} the maximum vehicle acceleration $(10 \,\mathrm{m/s^2})$
- \hat{j} the maximum vehicle jerk $(15 \,\mathrm{m/s^3})$

B. Consistency tests for object detection

$$\begin{array}{lll} \textbf{Edge} & \mathcal{O}_{\mathrm{R}} \rightarrow \mathcal{O}_{\mathrm{F}} \\ \textbf{Check} & U_R^F = \{o \in \mathcal{O}_{\mathrm{R}} \mid o \in \mathrm{FoV}(F) \wedge o \not \in \mathcal{O}_{\mathrm{F}} \} \\ & U_F^R = \{o \in \mathcal{O}_{\mathrm{F}} \mid o \in \mathrm{FoV}(R) \wedge o \not \in \mathcal{O}_{\mathrm{F}} \} \\ & U_F^R = \{o \in \mathcal{O}_{\mathrm{F}} \mid o \in \mathrm{FoV}(R) \wedge o \not \in \mathcal{O}_{\mathrm{R}} \} \\ & U = U_R^F \cup U_F^R \\ & \mathbf{return} \mid U \mid > 0 \text{, Severity}(U) \\ \\ \textbf{Edge} & \mathcal{O}_{\mathrm{F}} \rightarrow \mathcal{O}_{\mathrm{L}} \\ & U_F^L = \{o \in \mathcal{O}_{\mathrm{F}} \mid o \in \mathrm{FoV}(L) \wedge o \not \in \mathcal{O}_{\mathrm{L}} \} \\ & U_L^F = \{o \in \mathcal{O}_{\mathrm{L}} \mid o \in \mathrm{FoV}(F) \wedge o \not \in \mathcal{O}_{\mathrm{F}} \} \\ & U = U_F^L \cup U_L^F \\ & \mathbf{return} \mid U \mid > 0 \text{, Severity}(U) \\ \\ \textbf{Edge} & \mathcal{O}_{\mathrm{L}} \rightarrow \mathcal{O}_{\mathrm{R}} \\ & \mathcal{O}_{\mathrm{L}} \rightarrow \mathcal{O}_{\mathrm{R}} \\ & \mathcal{O}_{\mathrm{L}} = \{o \in \mathcal{O}_{\mathrm{L}} \mid o \in \mathrm{FoV}(L) \wedge o \not \in \mathcal{O}_{\mathrm{L}} \} \\ & U_R^L = \{o \in \mathcal{O}_{\mathrm{R}} \mid o \in \mathrm{FoV}(L) \wedge o \not \in \mathcal{O}_{\mathrm{L}} \} \\ & U = U_L^T \cup U_R^L \\ & \mathbf{return} \mid U \mid > 0 \text{, Severity}(U) \\ \end{array}$$

$$\begin{array}{ll} \textbf{Edge} & \mathcal{O}_{\mathrm{R}} \rightarrow \mathcal{O}_{\mathrm{C}} \\ \textbf{Check} & U_R^C = \{o \in \mathcal{O}_{\mathrm{R}} \mid o \in \mathrm{FoV}(C) \wedge o \not\in \mathcal{O}_{\mathrm{C}}\} \\ & U_C^R = \{o \in \mathcal{O}_{\mathrm{C}} \mid o \in \mathrm{FoV}(R) \wedge o \not\in \mathcal{O}_{\mathrm{R}}\} \\ & U = U_R^C \cup U_C^R \\ & \textbf{return} \ |U| > 0, \ \mathrm{Severity}(U) \end{array}$$

The set U_A^B represents the object detected by A that are in B field-of-view but we were unable to match (thus unmatched). The set $U=U_A^B\cup U_B^A$ represents the set object that one of the two sensors detected but the other did not. If the set U is non-empty, the consistency test reports a fault. The severity is computed as follow

$$\text{Severity}(U) = \begin{cases} \text{HIGH} & \text{if } \exists o \in U \text{ s.t. } o \in \text{CurrentLane} \\ \text{LOW} & \text{if } \exists o \in U \text{ s.t. } o \in \text{Map} \setminus \text{CurrentLane} \\ \text{NONE} & \text{otherwise} \end{cases}$$

As temporal consistency tests we use the same consistency tests we described above but with the proper change of inputs.

C. Consistency tests for vehicle localization

Given p_{t_1}, v_{t_1} and p_{t_2}, v_{t_2} the position and velocity estimated by the pose estimation module at the time t_1 and t_2 , and the set of acceleration measurements $a_i \dots a_j$ collected by the IMU in the interval $[t_1, t_2]$:

Given p_e , v_e the position and velocity estimated by the pose estimation module, and the position and velocity p_g , v_g measured by the GPS:

$$\begin{array}{ll} \textbf{Edge} & \text{POSE} \rightarrow \text{GPS} \\ \textbf{Check} & r = |p_e - p_g| \leq \tau_p \wedge |v_p - v_g| \leq \tau_v \wedge \dots \\ & v_e \leq \hat{v} \wedge v_g \leq \hat{v} \\ & \textbf{return} \ \neg r \end{array}$$

Given p_{t_1}, v_{t_1} and p_{t_2}, v_{t_2} the position and velocity measured by the GPS, and the set of acceleration measurements $a_i \dots a_j$ collected by the IMU in the interval $[t_1, t_2]$:

$$\begin{split} \textbf{Edge} & \quad \text{GPS} \rightarrow \text{IMU} \\ \textbf{Check} & \quad r = |p_{t_2} - p_{t_1}| \leq \hat{v}(t_{t_2} - t_1) \wedge \dots \\ & \quad |(v_{t_2} - v_{t_1}) - \int_{t_1}^{t_2} a(t) dt | \leq \tau_v \wedge \dots \\ & \quad v_{t_1} \leq \hat{v} \wedge v_{t_2} \leq \hat{v} \\ & \quad \textbf{return} \ \neg r \end{split}$$

Given p_{t_1}, v_{t_1} and p_{t_2}, v_{t_2} the position and velocity measured by the GPS at time t_1 and t_2 :

$$\begin{array}{ll} \textbf{Edge} & \text{GPS} \rightarrow \text{GPS} \\ \textbf{Check} & r = |p_{t_2} - p_{t_1}| \leq \tau_p \wedge |v_{t_2} - v_{t_1}| \leq \tau_v \wedge \dots \\ & v_{t_1} \leq \hat{v} \wedge v_{t_2} \leq \hat{v} \\ & \textbf{return} \ \neg r \\ \end{array}$$

Similarly, given p_{t_1}, v_{t_1} and p_{t_2}, v_{t_2} the position and velocity estimated by the pose estimation filter at time t_1 and t_2 :

 $\begin{array}{ll} \textbf{Edge} & \text{POSE} \rightarrow \text{POSE} \\ \textbf{Check} & r = |p_{t_2} - p_{t_1}| \leq \tau_p \wedge |v_{t_2} - v_{t_1}| \leq \tau_v \wedge \dots \\ & v_{t_1} \leq \hat{v} \wedge v_{t_2} \leq \hat{v} \\ & \textbf{return} \neg r \end{array}$

APPENDIX C EXTRA RESULTS





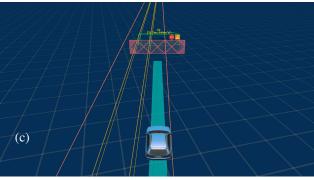
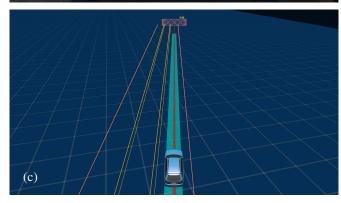




Fig. 1. **Simple scenario variation.** (a) Simulator, (b) Camera image, (c) World model (Apollo Auto Dreamview), (d) output of PerSyS. Apollo Auto sensor fusion correctly detects the overturned truck in the current lane as obstacle but the camera fails to do so, PerSyS is able to detect the inconsistency and report a high severity fault (red portion of the trajectory).







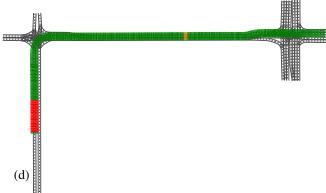


Fig. 2. **Night scenario variation.** (a) Simulator, (b) Camera image, (c) World model (Apollo Auto Dreamview), (d) output of PerSyS. Apollo Auto sensor fusion fails to detect the overturned truck in the current lane, PerSyS is able to detect the inconsistencies in the current lane between the radar-based obstacle detection and the other modules and report high severity faults (red portion of the trajectory).

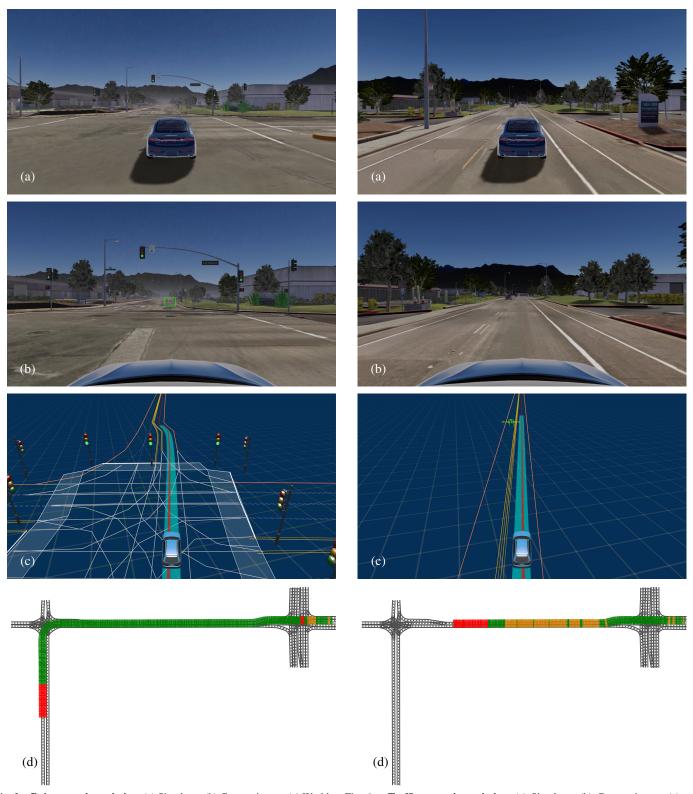
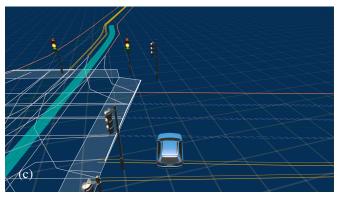


Fig. 3. Rain scenario variation. (a) Simulator, (b) Camera image, (c) World model (Apollo Auto Dreamview), (d) output of PerSyS. In adverse weather conditions the camera-based object detection wrongly detects an obstacle on the road floor, PerSyS detects the inconsistency and reports a high severity fault (red portion of the trajectory on the top-right). The red portion at the bottom-left corresponds to the high severity fault due to the overturned truck.

Fig. 4. **Traffic scenario variation.** (a) Simulator, (b) Camera image, (c) World model (Apollo Auto Dreamview), (d) output of PerSyS. Apollo Auto sensor fusion correctly detects the incoming car on the opposite lane but the camera fails to do so, as a result PerSyS reports low severity faults (yellow). In the final section of the trajectory PerSyS detects high severity faults (red) due to the overturned truck occupying the lane.







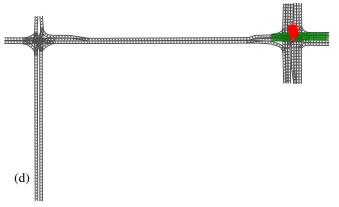


Fig. 5. **Spoiled GPS measurements.** (a) Simulator, (b) Camera image, (c) World model (Apollo Auto Dreamview), (d) output of PerSyS. Apollo Auto mislocalizes the AV due to spoiled GPS measurements, PerSyS detects inconsistencies between IMU and GPS and report a fault (red portion of the trajectory).