

## RESEARCH ARTICLE

# Evaluating Object (Mis)Detection From a Safety and Reliability Perspective: Discussion and Measures

ANDREA CECCARELLI<sup>1</sup> AND LEONARDO MONTECCHI<sup>2</sup>, (Member, IEEE)

<sup>1</sup>Department of Mathematics and Informatics, University of Florence, 50134 Florence, Italy

<sup>2</sup>Department of Computer Science, Norwegian University of Science and Technology (NTNU), 7034 Trondheim, Norway

Corresponding author: Andrea Ceccarelli (andrea.ceccarelli@unifi.it)

This work was supported in part by the European Commission (EC) under Grant H2020-MSCA-823788-ADVANCE; and in part by the National Recovery and Resilience Plan, Mission 4 Component 2-Investment 1.4-National Center for High Performance Computing (HPC), Big Data and Quantum Computing funded by the European Union-NextGeneration, under Grant EU-CUP B83C22002830001.


**ABSTRACT** We argue that object detectors in the safety critical domain should prioritize detection of objects that are most likely to interfere with the actions of the autonomous actor. Especially, this applies to objects that can impact the actor's safety and reliability. To quantify the impact of object (mis)detection on safety and reliability in the context of autonomous driving, we propose new object detection measures that reward the correct identification of objects that are most dangerous and most likely to affect driving decisions. To achieve this, we build an object criticality model to reward the detection of the objects based on proximity, orientation, and relative velocity with respect to the subject vehicle. Then, we apply our model on the recent autonomous driving dataset nuScenes, and we compare nine object detectors. Results show that, in several settings, object detectors that perform best according to the nuScenes ranking are not the preferable ones when the focus is shifted on safety and reliability.

**INDEX TERMS** Autonomous driving, object detection, safety, reliability.

## I. INTRODUCTION

The goal of object detection is to perceive and locate instances of semantic objects of a certain class [20]. A multitude of solutions have been proposed for 2D and 3D object detection, based on cameras and lidars [26], [44]. Object detection is fundamental in emerging safety-critical applications, and in particular it is a major pillar of autonomous driving applications [32].

To study object detection, new datasets are continuously proposed, for example in the autonomous driving domain KITTI [15], VOC [12], CityScapes [10], and more recently Waymo [35], nuScenes [5], and Level5 Lyft [19]. Object detectors are evaluated against datasets using widely acknowledged measures [30], [31], and well-defined routines [5], [10], [12], [15], [19], allowing a fair comparison.

The associate editor coordinating the review of this manuscript and approving it for publication was Divanilson Rodrigo Campelo .

Noteworthy, the Average Precision, first presented in [34], is currently deemed the most suitable measure to compute and rank the performance of object detectors.

However, we argue that current measures for object detection do not match the demands and peculiarities of autonomous vehicles and safety-critical systems in general, i.e., systems whose failure may lead to harmful consequences [3]. Evaluations based on Average Precision typically judge how well a detector detects objects, without discriminating based on the current position of these objects, and on their possibility to interfere with the subject in the considered scenario. To clarify, let us consider the typical modular pipeline for autonomous driving [17]: the subject vehicle is sensing the surroundings to perform object detection, and the output of the object detection is used for trajectory planning. Let us now consider two other vehicles in the sensed scenario, one directed straight towards the subject vehicle, in a colliding trajectory, and one headed away from the subject vehicle at

a higher speed. Clearly, for the safety of the driving task, it is critical to detect the first one, while detection of the second vehicle is not relevant at all. Unfortunately, this is not captured by the measures currently used in object detection, which consider both objects as equally relevant. Very practically, in a typical autonomous driving modular pipeline, it is first essential to detect all relevant objects, then these objects can be used for, e.g., trajectory planning. We argue that it is desirable the object detector does not fail to detect objects in colliding trajectory, otherwise also the output of the trajectory planner is compromised.

In this paper, we elaborate on how to measure the performance of object detectors in the safety-critical domain, with specific contextualization to the domain of autonomous driving, and we identify the need of an object criticality model and related measures. As key requirement, the desired measures should reward the detection of those objects that may interfere with the subject vehicle, and that are relevant for the safe and reliable execution of the driving task. Also, to be practically useful, the proposed measures have to be in a defined range, and be summarized by an overarching unifying measure. While autonomous driving is the most evident application domain, and it will be used as reference in the rest of this paper, our reasoning applies to any domain where reliability and safety of the object detection task are relevant for the success of the mission, for example in case of navigation and collision avoidance in drone systems [42].

More in detail, we propose a set of new measures, that we refer to as object criticality model. Such an object criticality model assigns a criticality score to each object, based on ground truth and estimated object distance, colliding trajectory, and time to collisions. Such criticality scores contribute to compute measures, named reliability-weighted precision and safety-weighted recall, that weight correct object detections and misdetections based on the impact on the safety and the reliability of the driving task. Last, a summarizing measure, named Critical Average Precision, allows ranking detectors according to such safety- and reliability-oriented measures.

The object criticality model and the related measures are exercised on the nuScenes dataset, with nine 3D-object detectors. We show that, under numerous settings, the ranking we obtain differs from the one achieved using the nuScenes evaluation library, which relies on traditional measures. Amongst implications, this result questions the usual approach to rate and select the most suitable object detector for the autonomous driving domain.

The rest of the paper is organized as follows. Section II presents basic notions and the related works. Section III shows the object criticality model and the measures we are introducing. Section IV describes the experiments based on nine object detectors and the nuScenes dataset. Section V illustrates the results, in which the object detectors are ranked according to our and traditional measures, and differences are discussed. Section VI concludes the paper.

## II. BACKGROUND AND RELATED WORKS

### A. OBJECT DETECTION AND ITS EVALUATION

We report the minimal set of notions on object detection that we require to present the choices made in our work.

To describe the spatial location and extent of a detectable object, in this paper for simplicity we only consider bounding boxes, although alternative approaches, e.g., [26], are applicable to our object criticality model as well.

Object detectors compute bounding boxes with an assigned confidence score. Then, a *detection threshold* is applied as a configuration parameter: all bounding boxes with a confidence score above the detection threshold are predictions. The classification of true positives (TPs), false positives (FPs), and false negatives (FNs), is based on some definition of distance between the predicted bounding boxes and the ground truth bounding boxes. In this paper, we use the distance between their center points [5]: a detected object is considered a TP if the distance between the ground truth bounding box and the detected bounding box is closer than a *distance limit*.

If there is no predicted bounding box that matches this criterion, then the object is not detected and it counts as an FN. Predicted bounding boxes that are farther than the distance limit from all ground truth bounding boxes are considered FPs. True negatives (TNs) are not taken into account, because there are infinite bounding boxes that should not be detected within any given image [30].

While there are several measures that can evaluate the performance of object detectors, the conventional approach to the evaluation of object detectors consists of measures that are derived from the count of TP, FP, and FN. These form the basis for our object criticality model defined in Section III, and they are briefly reviewed here [6], [30], [31]. *Precision*,  $P = TP/(TP + FP)$ , indicates how many of the selected items are relevant. If some non-relevant items are selected, this reduces precision. Precision is 1 if all the detected objects exist, and 0 in the opposite case. Conversely, *Recall*,  $R = TP/(TP + FN)$ , indicates how many of the existing relevant items are selected. If a detector has recall 1, it means it detected everything without any detection miss; in the opposite case, recall is 0. An object detector with high recall but low precision outputs many predictions, but most of them are incorrect; an object detector with high precision but low recall returns very few predictions, but most of them are correct.

Currently, the most frequently used summarizing measure is *Average Precision (AP)* [13], which summarizes the precision-recall curve as the weighted mean of precision scores achieved at different detection thresholds, using the increase in recall from the previous detection threshold as the weight. More precisely,  $AP = \sum_n (R_n - R_{n-1}) P_n$ , where  $P_n$  and  $R_n$  are the precision and recall at the  $n$ -th detection threshold. In this paper, in agreement with [5], we calculate AP only for recall and precision above or equal to 0.1: we remove cases in which recall or precision is less than 0.1 in

order to minimize the impact of noise commonly seen in regions with low precision or low recall.

### B. RELATED WORKS ON OBJECT DETECTION IN SAFETY-CRITICAL SYSTEMS

A safety-critical (computer) system is one whose malfunction could lead to unacceptable consequences, like harm to users or to the environment. A typical example is an autonomous vehicle, whose malfunction (of whatever cause) may lead to a collision. Instead, *reliability* describes the continuity of correct service, which can be temporarily disrupted, for example, to avoid situations that are potentially dangerous [3].

The inclusion of object detection tasks in safety-critical systems comes with a relevant set of renowned challenges, because of the many distinguishing aspects of the problem, its complexity, and also the variety of applications [1], [22], [38]. Considering object detectors, some incorrect predictions may lead to catastrophic consequences, and therefore have the maximum impact on safety, while others may have an irrelevant impact. Further, some false positives may cause an unnecessary interruption of the service, and therefore they impact the reliability. However, to evaluate object detectors, measures from Section II-A are typically used, without considering the different impact of each detection mistake. This also applies to the wide domain of autonomous driving, and it becomes evident when considering the measures used in object detection challenges for autonomous driving. For example, in challenges for KITTI [15], CityScapes [10], Waymo [35], or nuScenes [5], evaluation measures revolve around Average Precision and the concepts summarized in Section II-A.

Up to now, very few approaches have attempted to define safety or reliability measures for object detectors; to the best of our knowledge, the few works which targets a goal similar to ours are focusing on safety but leaving aside the reliability concern, and are [4], [28], [36], [39]. Noteworthy, they all appeared in very recent years, which underlines a recent understanding of the relevance of the subject, and they are all in the autonomous driving domain. The work in [4] ranks each object in three categories (imminent collision, potential collision, no collision), based on its collision risk. Instead, in [36] the authors define critical zones in which accurate perception is mandatory. On a similar position, the authors of [28] argue the relevance of identifying a distance up to which all pedestrians are detected. The closest approach to our work is [39], where the authors combine scores measuring detection quality, collision potential, and time needed to make the detection. This allows computing a safety score of a test scenario, in 5 classes from insufficient to excellent.

With respect to the reviewed works, the object criticality model we propose includes both safety and reliability issues of the driving task. This is important, because safety by itself (to detect everything which is potentially dangerous) can be enforced by low precision and high recall, i.e., low

false negatives at the cost of many false positives. Instead, by balancing both reliability and safety issues, our object criticality model provides a standalone evaluation of object detectors.

Other works address the problem of deep neural network uncertainty in autonomous driving, where the term uncertainty should be interpreted in the broad sense of how certain an object detector is about its predictions [14]. In general, these works aim to improve object detection, but they do not target the definition of specific measures. More specifically, the work in [14] argues that object detectors should also include prediction confidence, and it presents various methods to capture uncertainties in object detection for autonomous driving. Otherwise, object detectors can only tell the human drivers what they have seen, but not how certain they are about it. The work in [27] includes information on uncertainty sources (e.g., sensor noise), the work in [18] includes uncertainty when computing the bounding box regression loss, and the work in [21] considers both the noise inherent to the observations and the uncertainty that can be explained away given enough data. Last, despite not focusing on object detection, the work in [16] defines safety-oriented measures by proposing that predictions with a confidence score close to the detection threshold should be treated differently and more suspiciously. Finally, the work in [8] introduces the distinction of a critical area, which is the area nearby the vehicle where failed detection of an object may lead to immediate safety risks. The work acknowledges that the design of a driving application is focused on both i) guaranteeing safety in such critical area, and ii) guaranteeing high detection accuracy on the non-critical area (in order to have smooth driving). This observation leads the author to build different DNNs for the detection of objects in the two areas.

Still, the above works weight all the detected objects the same, i.e., when assessing the object detector, the usual binary (yes/no) counting of TPs, FPs, and FNs is performed. Instead, in our work we claim that i) object detectors should be evaluated depending on the ability to detect those objects that are most likely to affect the driving task, i.e., impact on safety and reliability, and ii) this can be realized by weighting the objects based on their criticality, and by building specific measures that consider such weights. Also, we remark that, in our object criticality model, measurement errors and uncertainty in the detection are inherently considered, when computing the scores assigned to each object, and when the predicted values are compared to the ground truth.

### III. OBJECT CRITICALITY MODEL

Our object criticality model is based on assigning a *criticality* value to each object in the scene, and then computing object detection measures that consider this criticality. The description of such model is independent of the sensors used to capture the scene (e.g., cameras or lidars) and of the type of objects.

## A. REQUIREMENTS AND ASSUMPTIONS

The application of the object criticality model requires i) a subject vehicle (named *ego* afterwards) that captures the scene with sensors as cameras and lidars, and ii) objects (other vehicles, pedestrians, etc.) that are within line-of-sight to *ego* and that are consequently captured by the sensors. This is the very typical situation of an autonomous vehicle that performs object detection.

We assume that the following ground truth information is available: i) 3D bounding boxes describing the size of the objects; ii) coordinates of *ego* and of the objects; iii) velocity of *ego* and of the objects. The most recent automotive datasets have very rich meta-data, typically including the above information; for example, in Section IV and in Section V we will use nuScenes [5], which satisfies our assumptions. Clearly, the ground truth is required only to evaluate the object detector, and not in the case of operation in a deployed setting.

Further, we assume that the object detector produces as output: i) the computed 3D bounding boxes, ii) the estimated distance of detected objects from *ego*, and iii) the estimated velocity of objects. In other words, the object detector is assumed to conflate detection, tracking, and dynamics: this is done in several 3D object detectors, which include the above estimates in their output. Noteworthy, these estimates are computed in the object detection challenges of the nuScenes community, which will be our reference for the experiments in Section IV and Section V.

For simplicity of the discussion, when computing coordinates of objects and their distance from *ego*, in this paper we consider only the  $(x, y)$  coordinates, i.e., we ignore the vertical dimension. In other words, while extending the object criticality model to the  $z$ -dimension is definitely possible, only at the cost of slightly more complex geometric computations, in the following we exclude the relative altitude of the objects and the *ego*. From the point of view of results, this is not an issue, because the dataset we use in this paper was collected on essentially flat lands. Also, note that ignoring possible vertical offsets of objects may only *reduce* their distance from *ego*, and it is therefore a worst-case approximation.

## B. STRUCTURE OF THE OBJECT CRITICALITY MODEL

We call *ego* the roving vehicle that mounts the sensors and collects data from the environment, and we call object *B* any other object. There are no restrictions on the type of objects, for example *B* can be a car, a pedestrian, a bike, etc. Note that for *ego* we only have ground truth values, i.e., the object detector does not predict its own velocity or position.

The construction of our object criticality model is organized in 3 steps, which are repeated for each object *B* within the line of sight of *ego*, and for both the ground truth values and the predicted values of *B*.

The first step (Section III-C) is the analysis of the collision scenario involving *B* and *ego*. In this step, we calculate indicators that will be later used to define the criticality

of *B*. In particular, we calculate i) the initial distance  $d$  between *ego* and *B*, ii) the closest distance  $r$  that *ego* and *B* would reach, and iii) the time  $\Delta t$  that *ego* and *B* require to reach such distance. These values are input to the following step, together with the current position and velocity of *ego* and *B*.

The second step (Section III-D) is the calculation of *criticality weights* that are assigned to each object *B*. These are  $\kappa_d$ ,  $\kappa_r$ , and  $\kappa_t$ , and they are based, respectively, on the three values calculated in the first step.

These weights indicate the relevance of *B* for the driving task: weights are higher if it is more likely that *B* may affect the behavior of *ego*. Such weights are used as rewards or penalties depending, respectively, on whether the object has been detected or missed.

The third step (Section III-E) exploits the assigned criticality to construct aggregate *safety and reliability measures* that allow comparing different detectors.

## C. ANALYTICAL CHARACTERIZATION OF THE COLLISION SCENARIO

We refer to Figure 1 for a visual representation of the collision scenario analyzed in this section.

We define  $ego = (ego_x, ego_y)$  the position of *ego*, and  $B = (B_x, B_y)$  the position of the object *B* in the captured scene. Further, we define, in vector form,  $v_{ego} = (v_{ego_x}, v_{ego_y})$  the velocity of *ego*, and  $v_B = (v_{B_x}, v_{B_y})$  the velocity of *B*. We compute the relative velocity of *B* with respect to *ego*, as  $v_{rel} = (v_{rel_x}, v_{rel_y}) = (v_{B_x} - v_{ego_x}, v_{B_y} - v_{ego_y})$ , that is, the vectorial difference of the velocity of *ego* and the velocity of *B*. This allows simplifying the subsequent calculations: we can consider *ego* as stationary, while *B* is moving with the velocity resulting from the difference between the two velocity vectors  $v_{ego}$  and  $v_B$ .

Then, we identify the shortest distance from *ego* at which object *B* will pass if both continue moving with the same velocity. This is the distance between *ego* and point  $C = (C_x, C_y)$ , with *C* being the point closest to *ego* on the trajectory of *B*. Point *C* can also be thought as the tangent point between the line representing the trajectory of *B* and a circle centered on *ego*.

Point  $C = (C_x, C_y)$  can be computed as the intersection of two lines, using basic Euclidean geometry. The line defining the direction of the relative movement of *B* is obtained from the general equation of a line, i.e.,  $y - y_0 = m(x - x_0)$ . We are looking for the line passing from point  $(B_x, B_y)$  and whose angular coefficient (i.e., orientation with respect to the  $x$  axis) is given by the ratio between the  $y$  and  $x$  components of the relative velocity  $v_{rel}$  (refer again to Figure 1).

The shortest distance between such line  $j$  and the position of *ego* lies on the line perpendicular to  $j$  passing from *ego*. By definition, a line perpendicular to one having coefficient  $m$  has coefficient  $m' = -1/m$ . Point  $C = (C_x, C_y)$  is then the point at which these two lines intersect, which is obtained by solving Section 1 below. Computations are

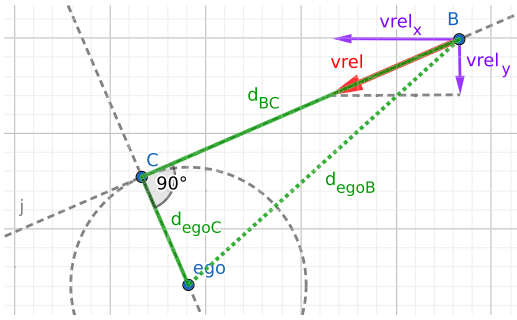


FIGURE 1. Geometrical representation of the main elements of our object criticality model.

omitted for brevity.

$$C_y = \frac{v_{rel_y}}{v_{rel_x}} (C_x - B_x) + B_y,$$

$$C_y = -\frac{v_{rel_x}}{v_{rel_y}} (C_x - ego_x) + ego_y, \quad (1)$$

Then, applying the Euclidean distance, we can easily compute: i) the distance  $d_{egoB}$  between *ego* and *B*, ii) the distance  $d_{egoC}$  between *ego* and *C*, and iii) the distance  $d_{BC}$  between *B* and *C*.

Assuming that both *ego* and *B* continue moving with the same velocity, the time  $\Delta t$  that *B* needs to reach the collision point *C* is then computed as the distance divided by the scalar speed of *B*, i.e.,  $\Delta t = d_{BC}/|v|$ , where  $|v| = \sqrt{v_x^2 + v_y^2}$ . We recall that *ego* is considered to be stationary, while *B* moves with a relative velocity obtained as the difference of the velocities of the two objects.

We note that including acceleration would better characterize objects' movement; however, since acceleration is quadratic with respect to space, any estimation error would be greatly amplified, introducing unnecessary noise in the final measures.

Finally, note that the object criticality model exhibits some corner cases, for example, when *ego* and *B* are moving at the exact same velocity,  $\Delta t$  is undefined. We treat these rare cases by skipping the object criticality model calculation and setting the criticality values to conservative fallback values.

### D. COMPUTATION OF CRITICALITY WEIGHTS

The collision scenario above is used to assign criticality to objects. The idea is inspired by reliability analysis [37], in which quantities like *reliability* (or *safety*) are defined in the interval [0, 1]. However, we do not propose probabilities.

Each object *B*, either identified by the object detector or ground truth, is assigned a *criticality weight*  $\kappa(B)$ . This weight is obtained by combining three criticality values  $\kappa_d(B)$ ,  $\kappa_r(B)$ , and  $\kappa_t(B)$ , as explained later. Note that for a given object *B*, its criticality  $\kappa(B)$  may be different if calculated with its predicted properties (e.g., position and velocity) or the ground truth ones. Furthermore, for some objects, we may have ground truth values only (FNs) or predicted values only (FPs). When needed, we indicate with  $\kappa'(B)$

the criticality weight computed with *predicted* properties of object *B*, as opposed to  $\kappa(B)$  that is calculated based on the ground truth.

#### 1) CRITICALITY SCORES

The *Distance Criticality*,  $\kappa_d(B)$ , is based on the distance  $d_{egoB}$  between *ego* and the object *B*. This score does not depend on velocity, but only on the position of objects in the scene. We want the score to be maximum when the distance from *ego* to *B* is zero, and then decrease to zero when reaching a maximum distance  $D_{max} > 0$ .

We compute the weight  $\kappa_d(B)$  as a second-degree equation (downward parabola) passing from points (0, 1) and ( $D_{max}$ , 0). That is, the maximum value is 1.0 when  $d_{egoB} = 0$  and it decreases as  $d_{egoB}$  increases, reaching 0 when  $d_{egoB} = D_{max}$ . The parabola shape allows the criticality to decrease non-linearly with respect to the distance: the decrease is slow for values close to zero (i.e., close to the vehicle), and it gets faster when approaching  $D_{max}$  (i.e., far from the vehicle). We also need to enforce that  $\kappa_d(B)$  is always in the interval [0, 1], and therefore the final equation is:

$$\kappa_d(B) = \max\left(0, -\frac{1}{Z^2}x^2 + 1\right), \quad x = d_{egoB}, \quad Z = D_{max}. \quad (2)$$

The *Collision Distance Criticality*,  $\kappa_r(B)$ , is based on the distance between *ego* and the potential collision point *C*. It is an indicator of how close to *ego* the object is likely to pass.  $\kappa_r(B)$  is calculated using the same rationale of  $\kappa_d(B)$  (Section 2), with  $x = d_{egoC}$  and  $Z = R_{max}$ , where  $R_{max} > 0$  is the maximum considered collision distance, beyond which the corresponding criticality is zero.

Similarly, the *Collision Time Criticality*,  $\kappa_t(B)$ , is based on the time  $\Delta t$  for *B* to reach the potential collision point. All the other things unchanged, this score depends on the (relative) velocity of the object *B* with respect to *ego*. This score is again calculated based on Section 2, with  $x = \Delta t$  and  $Z = T_{max}$ .

#### 2) OBJECT CRITICALITY

The final criticality  $\kappa(B)$  is obtained by the combination of the three criticality scores  $\kappa_d(B)$ ,  $\kappa_r(B)$ ,  $\kappa_t(B)$ . The resulting measure is defined following four requirements: i) it should range in the interval [0, 1]; ii) it should be 0 if all values are zero; iii) it should be 1 if at least one of the values is 1; and iv) it should increase if any of the three values increases.

Inspired again by classic reliability analysis [37], our final criticality weight is then computed as:

$$\kappa = 1 - (1 - \kappa_d) \cdot (1 - \kappa_r) \cdot (1 - \kappa_t). \quad (3)$$

The final criticality  $\kappa(B)$  is therefore a measure of: how much the object is close, how much it is likely to pass close in the near future, and how much time is available to react.

#### 3) CORNER CASES

The following corner cases are considered:

- When *ego* and *B* are moving at the exact same velocity (in both dimensions), the resulting relative velocity is

zero, and  $\Delta t$  cannot be computed. We solve this case by setting  $\kappa_r$  and  $\kappa_t$  to zero.

- The case in which only one component of the relative velocity is zero *does not* need to be treated differently. The resolution of Section III-C yields a form in which the denominator is the sum of squares of the two components of the velocity. The denominator is thus zero only when both components of the velocity are zero, which is already treated in the previous case.
- In the calculation of  $\Delta t$  we need to verify if the object B is actually moving towards point C, and not on the same line but in the opposite direction. In case B is moving in the opposite direction,  $\kappa_r$  and  $\kappa_t$  are again set to zero.
- In rare cases, where the collision point is particularly far away or the speed is particularly low, the calculation of  $\Delta t$  may generate an overflow or a not-a-number (NaN) value: in this case  $\kappa_t$  is set to 0.1. The rationale is to set it to a low value, but still greater than zero.
- The dataset may contain invalid values, or the detector may not be able to provide estimates. In particular, when we are not able to obtain the velocity of the object, we set  $\kappa_r$  and  $\kappa_t$  to 1 (their maximum value).

## E. SAFETY AND RELIABILITY-BASED MEASURES

We exploit the above criticality scores to remodel the traditional recall and precision measures, such that they are more oriented towards reflecting the safety and reliability offered by object detectors.

### 1) RELIABILITY

Reliability measures the *continuity of correct service*. [3]. For a reliable driving task, a good object detector should *not* predict false positives that correspond to dangerous situations, because they could lead to an interruption of the driving task. For example, false positives may cause an unnecessary brake; instead, the continuity of the driving mission may require considering some risks of collision as unavoidable. This clearly conflicts with safety (which aims to minimize risks), but it is widely accepted that safety and reliability have different goals [3] and may be conflicting requirements.

For this reason, we measure the reliability of the detection task through a revised definition of *precision*. The idea is that false positives are penalizing the continuity of the driving process, with a greater impact the closer they are, or are likely to be, to ego. We weight TPs and FPs according to the criticality  $\kappa(B)$  of the associated object B. In simpler words, when a non-existing object is detected, we do not add 1 to the count of FPs, but instead we add its criticality; the same applies to TPs.

For a correctly detected object we may use the criticality computed either using the ground truth ( $\kappa$ ) or the predicted values ( $\kappa'$ ): we use ground truth values at the numerator, and predicted values at the denominator. The idea is that the detector might detect a greater criticality (denominator) than what is actually present (numerator), which reduces reliability of

the driving task. Also, clearly we do not have ground truth values for FPs, because those objects do not exist.

We can then define the *reliability-weighted precision* as:

$$P_{\mathcal{R}} = \frac{\sum_{B \in TP^*} \kappa(B)}{\sum_{B \in TP^*} \kappa'(B) + \sum_{B \in FP^*} \kappa'(B)}, \quad (4)$$

where  $TP^*$  is the set of true positive objects, and  $FP^*$  is the set of false positive objects. Note that the  $P_{\mathcal{R}}$  may in principle raise above 1, in case the detected criticality is significantly lower than the ground truth. To be consistent with the classic definition of precision, we limit the maximum value of  $P_{\mathcal{R}}$  to 1.

### 2) SAFETY

Safety is instead the *absence of catastrophic consequences* [3]. To ensure safety, the object detector must detect as many as possible of the dangerous objects, even at the cost of raising some false alarms. A safety measure should then reflect how much of the existing criticality has been detected by the object detector. The proposed measure is adapted from the recall, using the ground truth values at the denominator and the detected values at the numerator. Clearly, we do not have predicted values for FNs, which are objects that have been missed. Therefore, we define the *safety-weighted recall* as:

$$R_{\mathcal{S}} = \frac{\sum_{B \in TP^*} \kappa'(B)}{\sum_{B \in TP^*} \kappa(B) + \sum_{B \in FN^*} \kappa(B)} \quad (5)$$

where  $TP^*$  is the set of true positive objects and  $FN^*$  is the set of false negative objects. Also for  $R_{\mathcal{S}}$  we limit its maximum value to 1.

### 3) CRITICAL AVERAGE PRECISION

The proposed criticality values depend on three parameters, namely  $D_{max}$ ,  $R_{max}$ , and  $T_{max}$ . We can compute  $P_{\mathcal{R}}$  and  $R_{\mathcal{S}}$  for different values of these parameters, to understand their evolution when different subsets of objects are considered. In analogy to the precision-recall curve (see Section II-A), this allows computing several  $P_{\mathcal{R}}-R_{\mathcal{S}}$  curves, one for each combination of values ( $D_{max}$ ,  $R_{max}$ ,  $T_{max}$ ); consequently, we can compute the Critical Average Precision  $AP_{crit}$  from each of the  $P_{\mathcal{R}}-R_{\mathcal{S}}$  curves, based on our definitions of  $P_{\mathcal{R}}$  and  $R_{\mathcal{S}}$ .

Depending on the driving scenario and the intended system in which the object detector is deployed, different values of  $D_{max}$ ,  $R_{max}$ , and  $T_{max}$  may be favored. For example, an object detector which is very good on  $P_{\mathcal{R}}$  could be safely used on a highway under low traffic conditions; but if it is not good on  $R_{\mathcal{S}}$ , it should not be used in an urban scenario, where cars may approach from different directions at essentially any angle.

## IV. CASE STUDY ON THE nuScenes DATASET

### A. DATASETS AND SELECTED OBJECT DETECTORS

To exercise the object criticality model, we choose the nuScenes dataset for the following reasons: i) it is very recent and extensive, forged with the latest sensor technology; ii) very recent object detectors are available; iii) it includes all

the necessary information to apply the object criticality model presented in Section III.

NuScenes [5] is a recent large-scale dataset for autonomous driving that reports scenes collected from a vehicle. The dataset comprises 1000 scenes, each being 20 seconds long and fully annotated with 3D bounding boxes. *Keyframes* are sampled every 0.5 seconds; five intermediate frames are collected between keyframes.

Following common practices in datasets of this kind [15], [35], nuScenes defines an object detection task and proposes related measures to officially rank object detectors on its website. The detection task in nuScenes consists in predicting the objects at each keyframe time  $t$ , using sensors data collected between  $(t-0.5, t]$  seconds (five intermediate frames). Detectable objects are all objects within 50 meters from ego and with line of sight. For each object, ground truth 3D bounding boxes, attributes (e.g., orientation), and velocities are provided. A detection is successful if the distance between the centers of the predicted and ground-truth bounding boxes is less than a distance limit  $l$ ; four different values of  $l$  are considered, which are  $l \in \{0.5, 1, 2, 4\}$  meters. For brevity of the discussion, the only objects we consider are cars.

We select nine 3D object detectors from the zoo of mmdetection3d [9], an open-source object detection toolbox based on PyTorch for 3D detection. We present the object detectors below; each detector is matched to an acronym to easily distinguish it in the rest of the paper.

*FCOS* [40] and its evolution *PGD* [41] use visual cameras only. The backbone is a pretrained ResNet101 with deformable convolutions [11]. The neck is the Feature Pyramid Network (FPN, [24]), which generates a pyramid of feature maps. The head that produces final predictions (deciding on object class, location, etc.) relies on an approach similar to RetinaNet [25], which applies shared heads to operate detection of multiple targets. PGD head also includes a branch to improve the estimation of distance depth.

The other seven object detectors (see Table 1) process lidar's pointcloud and they are based on the Pointpillars [23] network. Pointpillars is well-known both for its speed and its accuracy. It exploits an encoder that learns features on pillars (vertical columns) of the point cloud to predict 3D oriented bounding boxes for objects. The Pointpillars network consists of three main stages: i) a feature encoder network that converts a point cloud to a structured representation, namely a sparse pseudoimage; ii) a 2D convolutional backbone to process the pseudo-image into high-level representation, extracting the features map upon which the rest of the network is used; and iii) a detection head that detects and regresses 3D bounding boxes. We consider seven alternatives based on Pointpillars; essentially, they use the pillar-based method from [23] to convert the point cloud into a sparse pseudoimage, and differentiate from [23] by applying different backbones, and optionally changing the necks and heads.

**TABLE 1. The seven lidar-based object detectors in use.**

Acronym	Short description
FPN	Backbone is FPN.
REG400	Backbone is the REGNETX-400MF from [33].
REGSEC	Similar to REG400, but includes the neck SECOND [43].
REG1.6	Backbone is the REGNETX-1.6GF DNN from [33].
SEC	Backbone is FPN, neck is SECOND [43].
SSN	As in SEC above, but it adds the shape-aware grouping heads from [45].
SSNREG	As in SSN, but the backbone is REGNETX-400MF from [33].

## B. IMPLEMENTATION OF THE OBJECT CRITICALITY MODEL

We execute all the object detectors on the nuScenes validation set [5], which consists of 150 frame sequences of 20 seconds each, and achieved the exact same results of their authors reported at [9]. This confirms that our setup of mmdetection3d is correct.

The implementation of our object criticality model exploits the development kit of nuScenes, which is available with open-source license. For example, the ranking of object detectors available at the nuScenes website [29] is computed using the code of this library, but on a different test set, whose ground truth information is not released to the public. We extended the development kit, to have it compute the measures from our object criticality model alongside the usual measures of the nuScenes object detection challenge. We compute and plot the analogous of the precision-recall curve, but with our criticality-oriented measures  $\mathcal{P}_{\mathcal{R}}$  and  $\mathcal{R}_{\mathcal{S}}$ . The resulting library is available at [7]. Its usage is straightforward: it is sufficient to have a working installation of nuScenes-dev, and replace with the files in [7] the corresponding files of the nuScenes-dev installation. Then, the set of results will appear enriched with our measures. Therefore, any object detector whose output is compatible with nuScenes can be also evaluated using our library. The library is released open source on [7], including tutorials that reproduce the experiments described in this paper. We used the nuScene development kit v1.1.2, and we tested for compatibility up to 1.1.7. The release at [7] includes a usage example, which allows repeating our experiments from the execution of the mmdetection3d object detectors to the computation of results.

## V. EXPERIMENTS AND RESULTS

We execute the 9 object detectors on the dataset previously described. We compute  $AP_{crit}$ ,  $\mathcal{P}_{\mathcal{R}}$  and  $\mathcal{R}_{\mathcal{S}}$  for different values of  $D_{max}$ ,  $R_{max}$ , and  $T_{max}$ . More specifically, we consider several configurations  $(D_{max}, R_{max}, T_{max})$ , with  $D_{max} \in \{5, 10, \dots, 50\}$  meters,  $R_{max} \in \{5, 10, \dots, 50\}$  meters, and  $T_{max} \in \{2, 4, \dots, 30\}$  seconds. Since distance is measured starting from the center of ego, a distance of 5 meters includes only vehicles very close to ego; 50 meters instead is the maximum distance from ego that is considered in the nuScenes object detection challenge, where objects farther than 50 meters from ego are ignored. Overall, this

**TABLE 2.**  $AP$  and  $AP_{crit}$  of car detection, for nine object detectors and  $l \in \{0.5, 1, 2, 4\}$ , ordered by  $AP$ .  $AP_{crit}$  is computed with  $(D_{max}, R_{max}, T_{max})$  amongst the configurations that reported the highest differences between  $AP_{crit}$  and  $AP$  ranking. Ranking differences are in bold.

(a) $l = 0.5, (20, 20, 8)$			(b) $l = 1, (20, 25, 10)$			(c) $l = 2, (10, 35, 8)$			(d) $l = 4, (20, 25, 4)$		
Detector	$AP$	$AP_{crit}$	Detector	$AP$	$AP_{CRIT}$	Detector	$AP$	$AP_{CRIT}$	Detector	$AP$	$AP_{CRIT}$
FCOS	0.118	0.1995	FCOS	0.372	0.4857	FCOS	0.655	0.7061	FCOS	0.804	0.8695
PGD	0.172	0.2711	PGD	0.450	0.5658	PGD	0.703	0.7382	PGD	0.828	0.8860
SEC	0.677	0.7534	SEC	0.796	0.8486	SEC	0.833	0.8611	SEC	0.852	0.9057
FPN	0.682	<b>0.7618</b>	FPN	0.814	<b>0.8631</b>	FPN	0.857	<b>0.8748</b>	FPN	0.872	<b>0.9202</b>
REG400	0.690	<b>0.7616</b>	SSN	0.818	<b>0.8628</b>	SSN	0.860	<b>0.8691</b>	REGSEC	0.875	<b>0.9205</b>
SSN	0.696	0.7717	REGSEC	0.825	<b>0.8728</b>	REGSEC	0.863	<b>0.8838</b>	SSN	0.876	<b>0.9200</b>
REGSEC	0.713	0.7773	REG400	0.827	<b>0.8703</b>	REG400	0.870	<b>0.8839</b>	REG400	0.884	<b>0.9267</b>
SSNREG	0.717	<b>0.7903</b>	SSNREG	0.835	<b>0.8806</b>	SSNREG	0.872	<b>0.8912</b>	SSNREG	0.886	<b>0.9293</b>
REG1.6	0.722	<b>0.7895</b>	REG1.6	0.837	<b>0.8769</b>	REG1.6	0.874	<b>0.8837</b>	REG1.6	0.889	<b>0.9264</b>

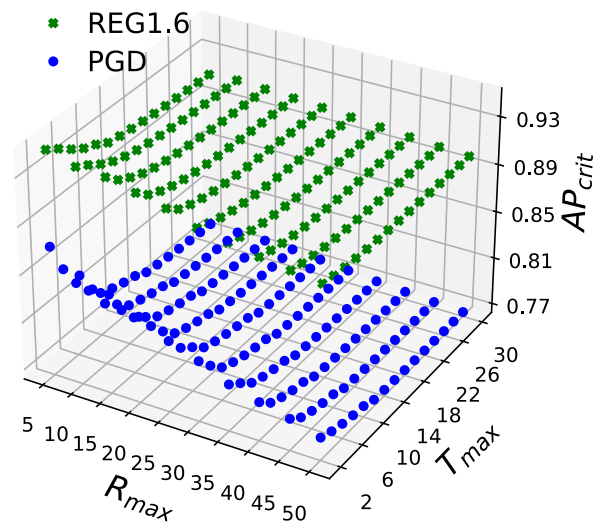
leads to 1500 configurations  $(D_{max}, R_{max}, T_{max})$ , repeated for each object detector.

**A.  $AP_{crit}$  AND RANKING OF OBJECT DETECTORS**

First, we calculate the rankings of detectors based on  $AP_{crit}$  for all the 1500 configurations  $(D_{max}, R_{max}, T_{max})$ . Many of them produced a different ranking with respect to the one based on  $AP$ . For example, consider  $l = 0.5$  and  $l = 4$ . When  $l = 0.5$ , the ranking calculated with  $AP_{crit}$  does not match the  $AP$  ranking for 567 out of 1500 configurations; for each of these 567 configurations, the difference with respect to the  $AP$  ranking is 2 or 4 positions. The whole set of object detectors may change position with respect to the  $AP$  ranking, with the exception of the detector in the 7<sup>th</sup> and 8<sup>th</sup> positions which are always PGD and FCOS, respectively. For  $l = 4$ , the ranking changes in 1425 out of 1500 configurations, and all the object detectors may change position, including FCOS performing better than PGD. In Table 2, we compare the  $AP$  and  $AP_{crit}$  ranking of the nine object detectors, for exemplary configurations  $(D_{max}, R_{max}, T_{max})$ . Noticeably, the object detector with the highest  $AP$ , REG1.6, is outperformed by SSNREG and also others when we consider  $AP_{crit}$ .


To explore trends of  $AP_{crit}$ , we select representative examples. In Figure 2 we show the  $AP_{crit}$  values of object detectors REG1.6 ( $AP = 0.874$ ) and PGD ( $AP = 0.703$ ) with  $l = 2$  and when  $D_{max} = 25$ , for different  $R_{max}, T_{max}$ . The  $AP_{crit}$  of REG1.6 and PGD is higher than the respective  $AP$ s under the considered configurations. In fact, setting  $D_{max} = 25$  reduces the impact of objects farther than 25 meters, which are a significant contribution to misdetections.

Next, we pick the object detector REG1.6 with  $l = 2.0$ . In Figure 3 we show the  $AP_{crit}$  when  $R_{max} = 20$ ; the figure clearly shows how the highest  $AP_{crit}$  values are achieved when  $D_{max}$  is set in the range [20, 30]. This is possibly due to the fact that setting  $D_{max}$  very low excludes a lot of “easy” (i.e., close) objects from the relevant ones, thus deteriorating  $AP_{crit}$ . Conversely, when  $D_{max}$  becomes much greater than  $R_{max}$ , a lot of distant but not relevant objects are included, which are unlikely to reach a collision point closer than  $R_{max}$ . In the lower part of the  $z$  axis,  $AP = 0.874$  is represented as a flat grey surface in the figure.



**FIGURE 2.**  $AP_{crit}$  measured on REG1.6 and PGD with  $l = 2$  and  $D_{max} = 25$ .

Figure 3 shows that  $AP_{crit}$  is in general higher than  $AP$ . This is expected, because the  $AP_{crit}$  gives less weight to objects that are harder to detect, e.g., those at a farther distance from ego.



**FIGURE 3.**  $AP_{crit}$  measured on REG1.6 with  $l = 4.0$  and  $R_{max} = 20$ , for the different  $D_{max}$  and  $T_{max}$ .



In general, higher values of  $AP_{crit}$  are achieved with low values of  $R_{max}$  and  $T_{max}$ : for both REG1.6 and PGD, the maximum  $AP_{crit}$  values are obtained with  $(D_{max}, R_{max}, T_{max}) = (25, 5, 2)$ . Intuitively, low  $R_{max}$  and  $T_{max}$  reduce the number of vehicles to be considered in our analysis: only those that are really critical for the detection are included. Analogous observations can be derived with the other configurations and object detectors.

We remark that, while studies like Figure 2 and Figure 3 are effective to explain the proposed  $AP_{crit}$  measure, the most suitable configuration  $(D_{max}, R_{max}, T_{max})$  should be decided based on the requirements of the target application, and then the object detector with the highest  $AP_{crit}$  for such configuration should be selected.

**B. TRADEOFF BETWEEN  $P_{\mathcal{R}}$  AND  $R_S$**

To discuss the relations between  $P_{\mathcal{R}}$  and  $R_S$ , we rely on Figure 4 and Figure 5, where we use SSNREG with  $l \in \{1, 4\}$ . We compute  $P_{\mathcal{R}}$  and  $P$  for, respectively,  $R_S$  and  $R$  at steps of 0.01, starting from 0.85. Red crosses represent precision-recall pairs  $(P, R)$ . Black dots represent  $(P_{\mathcal{R}}, R_S)$  pairs; these are computed for each configuration  $(D_{max}, R_{max}, T_{max})$ , thus yielding 1500 black dots for each  $R_S$  value. The large blue dots are the  $(P_{\mathcal{R}}, R_S)$  values achieved using SSNREG with the configuration from Table 2b and Table 2d, while the green triangle are the configuration leading to the highest  $AP_{crit}$ , which is  $(25, 5, 2)$ .




FIGURE 4.  $P_{\mathcal{R}}, R_S, P$  and  $R$  for SSNREG when  $R_S \geq 0.85$  and  $R \geq 0.85$ , with  $l = 1$ .

We investigate the relations between  $P_{\mathcal{R}}$  and  $R_S$  for high values of  $R_S$  (safety-weighted recall), which are of particular interest in the reference domain of this work. This way we can study the  $P_{\mathcal{R}}$  (reliability-weighted precision) that we achieve when safety is enforced thanks to a high  $R_S$ . This corresponds to answering the question “given a safety target on the detection, what is the possibility of driving the car with good mission reliability, i.e., without being forced to interrupt the driving continuously because of false positives?”. Of course, the safest condition would be  $R_S = 1$ , but  $P_{\mathcal{R}}$  is typically 0 in

such cases; still, a very high  $R_S$  is necessary to enforce safety of the detection.




FIGURE 5.  $P_{\mathcal{R}}, R_S, P$  and  $R$  for SSNREG when  $R_S \geq 0.85$  and  $R \geq 0.85$ , with  $l = 4$ .

When the recall  $R$  increases, the precision  $P$  quickly drops to 0. SSNREG can offer a high recall, i.e., a high ability to detect all the objects, only at the cost of many false positives: this is clearly of little or no use in practice. Instead, if we restrict the scope of the object detector thanks to our object criticality model, we reach different conclusions. For example, consider again the case  $l = 1$  (Table 2b). Even with  $R_S \geq 0.9$ , there are some configurations in which  $P_{\mathcal{R}} > 0.8$ , which is clearly a much more comforting result, showing confidence in the detection at least to some extent.

On the other hand, the best-performing triples, represented with the green triangles in Table 2b, may be not practical, because it is computed applying small spatial and temporal distances of the objects from ego. Summarizing, our conclusion on SSNREG can be very different from those we achieve using  $P$  and  $R$ , when we apply the criteria of  $R_S$  and  $P_{\mathcal{R}}$ .

**C. EXPLANATION OF DISTANCE CRITICALITY  $\kappa(B)$**

The objective of this analysis is to explain the inner details of the object criticality model, even if  $P_{\mathcal{R}}, R_S$ , and  $AP_{crit}$  are sufficient to describe the performance of the object detection. We rely on bird-views from selected frames of nuScene to explain how our object criticality model works, in a very practical way, for the computation of  $\kappa(B)$ . We consider PGD and SEC object detectors, but all nine detectors lead to similar conclusions.

Figures from Figure 6 to Figure 9 are extracted relying on the nuscene-dev kit 1.1.2, properly modified to visualize values from  $\kappa(B), \kappa_d(B), \kappa_r(B)$ , and  $\kappa_t(B)$ . The axes represent distances, in meters. The ego is always located in the center at the  $(0, 0)$  coordinates and is oriented along the y-axis (heading towards the top). The other vehicles are represented as rectangles, and the front side is indicated by a small segment. The ground truth (real position and orientation of cars) is in green, while the detected cars are in blue. In the ideal

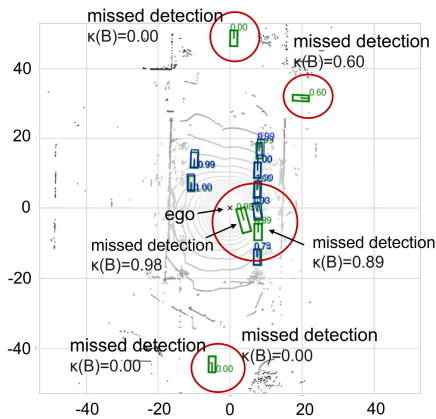


FIGURE 6.  $\kappa(B) = 0.98$  and  $\kappa(B) = 0.89$  for two dangerous missed detections from PGD.

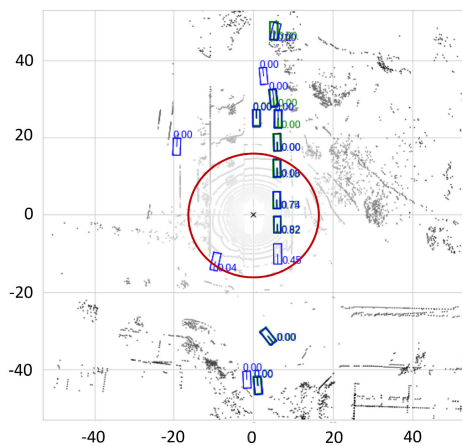


FIGURE 7.  $\kappa_d(B)$  computed for SEC with  $D_{max} = 15$ ,  $R_{max} = 20$ ,  $T_{max} = 10$ .

case of a perfect object detectors, blue and green rectangles would overlap. Both ground truths and detected vehicles have associated a value, which is either  $\kappa(B)$ ,  $\kappa_d(B)$ ,  $\kappa_r(B)$ , and  $\kappa_t(B)$  depending on the figure. We add text labels and red circles to improve readability.

We first consider PGD with  $D_{max} = 30$ ,  $R_{max} = 20$ ,  $T_{max} = 8.0$ . Very intuitively, this setting says that it is critical to detect vehicles that are within 30 meters, and/or that are in colliding trajectories within 20 meters in the next 8 seconds.

We start from Figure 6. A car is very close to ego, but it is not detected: it has been assigned  $\kappa(B) = 0.98$ . This car is located at the center of the diagram, and it is circled in red. Another one is very close, but in “a less dangerous” situation:  $\kappa(B) = 0.89$ . A third one is within  $D_{max} = 30$  meters, but headed in a different direction, so it gets a mild criticality score  $\kappa(B) = 0.60$ . Instead, there are other less critical missed detections in the upper and lower parts of the image. These are farther than  $D_{max} = 30$  meters, and are headed in non-colliding trajectories: these are irrelevant, so they are worth  $\kappa_d(B) = 0.00$ .

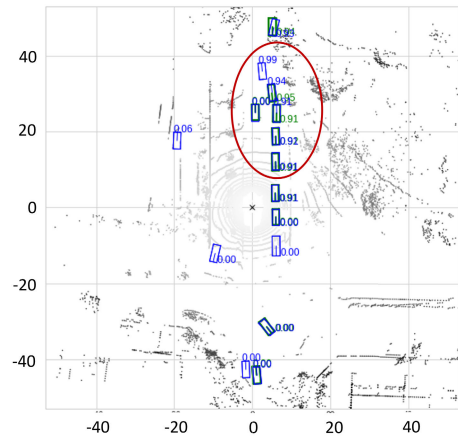


FIGURE 8.  $\kappa_r(B)$  computed for SEC with  $D_{max} = 15$ ,  $R_{max} = 20$ ,  $T_{max} = 10$ . The red circle is an area of approximately  $R_{max} = 20$  meters from ego.

Next, we explore the contribution of the distance criticality  $\kappa_d(B)$ . We consider SEC with  $D_{max} = 15$ ,  $R_{max} = 20$ ,  $T_{max} = 10$ . Figure 7 shows that cars farther than 15 meters from ego are assigned  $\kappa_d(B) = 0$ ; the closest to ego, the higher the  $\kappa_d(B)$  values. The red circle is approximately 15 meters radius: vehicles outside the circle have  $\kappa_d(B) = 0$ .

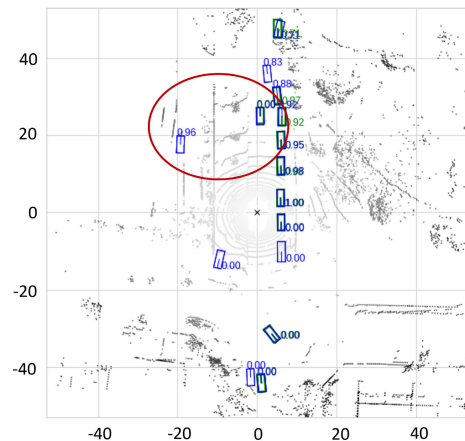


FIGURE 9.  $\kappa_t(B)$  computed for SEC with  $D = 15$ ,  $l = 20$ ,  $T = 10$ . Objects which may collide with ego within 10 seconds have assigned  $\kappa_t(B) \geq 0$ .

Figure 8 shows the values of  $\kappa_r(B)$  for the same scene and settings of Figure 7. Cars with a trajectory passing closer to ego than  $R_{max} = 20$  are assigned  $\kappa_r(B) > 0$ ; the closest to ego, the higher the values, meaning that those passing very close are at risk of collision. Note that objects can be assigned a high  $\kappa_r(B)$  value even if they appear to actually collide with ego. It is the case of the line of cars in the upper right part of Figure 8, in the red circle: even if they are proceeding straight, they are passing very close on the right of ego, and thus the indicator of potential collision is close to 1.0.

Similarly, Figure 9 shows the values of  $\kappa_t(B)$ . Cars which may enter in a collision within  $T = 10$  seconds are assigned

$\kappa_t(B) > 0$ . The velocity of ego and each car is a determining factor to assign the criticality  $\kappa_t(B)$ : vehicles relatively close and in colliding trajectory may also have  $\kappa_t(B) = 0$  values if they are not expected to reach the collision point within  $T = 10$ . In the red circle, there are opposite examples, of detected cars in colliding trajectories with ego but with  $\kappa_t(B) = 0$  and  $\kappa_t(B) = 0.96$  (note that this last one is a false positive).

## VI. CONCLUSION AND FUTURE WORKS

We argue that the most used measures for object detection do not match the demands and peculiarities of a safety-critical system. Within the autonomous driving domain, currently adopted measures typically describe how good an object detector is at detecting *all* the objects on the scene, while instead, for the purpose of an autonomous driving system, we are interested in detecting all the objects that will likely interfere with the driving task of the vehicle.

To this end, we show that the state-of-the-art evaluation of object detectors does not consider the possible role of the objects in a specific scene, and in particular with respect to the driving task of the vehicle performing the detection.

Consequently, we propose novel measures that take into account the concepts of safety (detection of dangerous objects, which require immediate reaction, should be prioritized) and reliability (mis-detections should not severely disrupt the continuity of the driving task). We build and exercise an object criticality model that performs a rating of the objects, based on the distance from the subject vehicle, the possible colliding trajectory, and the expected time to collision. Amongst the main results, we show that our judgment on the performance of object detectors may be very different when we consider the detection of i) everything on the scene (as it is usually done), or ii) only the relevant items. Depending on which of the two cases is of interest, we may end up choosing different object detectors. Further, we show that object detectors with high performance under case i) can be less competitive in case ii), and vice-versa.

Last, an important implication of our object criticality model is that, when safety and reliability issues are considered, the selection of the most suitable object detector strictly depends on its desired use, i.e., on the requirements of the target application. Starting from application requirements, the desired configuration of our object criticality model is identified, measures are computed, and the most suitable object detector is selected.

Noteworthy, our analysis is not meant to prove that the evaluated object detectors are safe and reliable. Rather, it shows how the object criticality model allows establishing sound parameters that can be used to build, assess and tune object detectors for their application in safety-critical domains.

We remark that object detection in complex scenarios is still an open research topic that makes improvements every year [2], with new detectors that are proposed continuously;

however, defining new object detectors, or assessing the most up-to-date object detectors, is beyond the scope of this paper.

As future work, we are currently working towards training an object detector whose goal is to maximise  $AP_{crit}$ . More precisely, the objective is to train to maximize a specific configuration of  $R_S$  and  $P_S$ , rather than  $R$  and  $P$ . Intuitively, the object detector is intended to reward the detection of objects that are relevant (close and in colliding trajectories), and it is expected instead to be far less effective in the detection of objects that are not relevant for the driving tasks and that do not interfere with the elaboration of the trajectory of ego. Practically, this can be realized by a proper training phase, where the usual loss measurement approach is modified according to the principles and measures established in this work.

## REFERENCES

- [1] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in AI safety," 2016, *arXiv:1606.06565*.
- [2] E. Arnold, O. Y. Al-Jarrah, M. Dianati, S. Fallah, D. Oxtoby, and A. Mouzakitis, "A survey on 3D object detection methods for autonomous driving applications," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 10, pp. 3782–3795, Oct. 2019.
- [3] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Trans. Depend. Secure Comput.*, vol. 1, no. 1, pp. 11–33, Jan. 2004.
- [4] A. Bansal, J. Singh, M. Verucchi, M. Caccamo, and L. Sha, "Risk ranked recall: Collision safety metric for object detection systems in autonomous vehicles," in *Proc. 10th Medit. Conf. Embedded Comput. (MECO)*, Jun. 2021, pp. 1–4.
- [5] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "NuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11621–11631.
- [6] G. O. Campos, A. Zimek, J. Sander, R. J. Campello, B. Mícenková, E. Schubert, I. Assent, and M. E. Houle, "On the evaluation of outlier detection: Measures, datasets, and an empirical study continued," in *Proc. Lernen, Wissen, Daten, Analysen 2016*, 2016, pp. 1–43.
- [7] A. Ceccarelli and L. Montecchi. *Library and Examples Used in This Paper*. Accessed: May 8, 2023. [Online]. Available: [https://github.com/AndreaCeccarelli/metrics\\_model](https://github.com/AndreaCeccarelli/metrics_model)
- [8] C.-H. Cheng, "Safety-aware hardening of 3D object detection neural network systems," in *Proc. Int. Conf. Comput. Saf., Rel., Secur.* Cham, Switzerland: Springer, 2020, pp. 213–227.
- [9] Contributors. (2020). *MMDetection3D: OpenMMLab Next-Generation Platform for General 3D Object Detection*. [Online]. Available: <https://github.com/open-mmlab/mmdetection3d>
- [10] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223.
- [11] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, "Deformable convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 764–773.
- [12] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes challenge: A retrospective," *Int. J. Comput. Vis.*, vol. 111, no. 1, pp. 98–136, Jan. 2015.
- [13] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, Jun. 2010.
- [14] D. Feng, L. Rosenbaum, and K. Dietmayer, "Towards safe autonomous driving: Capture uncertainty in the deep neural network for LiDAR 3D vehicle detection," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 3266–3273.
- [15] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, Sep. 2013.

- [16] M. Gharib, T. Zoppi, and A. Bondavalli, "Understanding the properness of incorporating machine learning algorithms in safety-critical systems," in *Proc. 36th Annu. ACM Symp. Appl. Comput.*, Mar. 2021, pp. 232–234.
- [17] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu, "A survey of deep learning techniques for autonomous driving," *J. Field Robot.*, vol. 37, no. 3, pp. 362–386, Apr. 2020.
- [18] Y. He, C. Zhu, J. Wang, M. Savvides, and X. Zhang, "Bounding box regression with uncertainty for accurate object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2888–2897.
- [19] J. Houston, G. Zuidhof, L. Bergamini, Y. Ye, L. Chen, A. Jain, S. Omari, V. Igloukov, and P. Ondruska, "One thousand and one hours: Self-driving motion prediction dataset," in *Proc. Conf. Robot Learn.*, vol. 155, J. Kober, F. Ramos, and C. Tomlin, Eds., Nov. 2021, pp. 409–418.
- [20] L. Jiao, "A survey of deep learning-based object detection," *IEEE Access*, vol. 7, pp. 128837–128868, 2019.
- [21] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" 2017, *arXiv:1703.04977*.
- [22] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," *IEEE Intell. Transp. Syst. Mag.*, vol. 9, no. 1, pp. 90–96, Spring 2017.
- [23] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "PointPillars: Fast encoders for object detection from point clouds," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 12697–12705.
- [24] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2117–2125.
- [25] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2980–2988.
- [26] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, and M. Pietikäinen, "Deep learning for generic object detection: A survey," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 261–318, Feb. 2020.
- [27] A. Loquercio, M. Segu, and D. Scaramuzza, "A general framework for uncertainty estimation in deep learning," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 3153–3160, Apr. 2020.
- [28] M. Lyssenko, C. Gladisch, C. Heinzemann, M. Woehle, and R. Triebel, "From evaluation to verification: Towards task-oriented relevance metrics for pedestrian detection in safety-critical domains," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2021, pp. 38–45.
- [29] nuScenes. (2022). *nuScenes Web-Site*. [Online]. Available: <https://www.nuscenes.org/>
- [30] R. Padilla, S. L. Netto, and E. A. B. da Silva, "A survey on performance metrics for object-detection algorithms," in *Proc. Int. Conf. Syst., Signals Image Process. (IWSSIP)*, Jul. 2020, pp. 237–242.
- [31] D. M. W. Powers, "Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation," 2020, *arXiv:2010.16061*.
- [32] C. Prenebida, G. Melotti, and A. Asvadi, "RGB-D object classification for autonomous driving perception," in *RGB-D Image Analysis and Processing*. Cham, Switzerland: Springer, 2019, pp. 377–395.
- [33] I. Radosavovic, R. P. Kosaraju, R. Girshick, K. He, and P. Dollar, "Designing network design spaces," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10425–10433.
- [34] G. Salton and M. J. McGill, *Introduction to Modern Information Retrieval*. New York, NY, USA: McGraw-Hill, 1983.
- [35] P. Sun, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 2443–2451.
- [36] S. Topan, K. Leung, Y. Chen, P. Tupekar, E. Schmerling, J. Nilsson, M. Cox, and M. Pavone, "Interaction-dynamics-aware perception zones for obstacle detection safety evaluation," 2022, *arXiv:2206.12471*.
- [37] S. K. Trivedi, *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*. Upper Saddle River, NJ, USA: Prentice-Hall, 1982.
- [38] K. R. Varshney and H. Alemzadeh, "On the safety of machine learning: Cyber-physical systems, decision sciences, and data products," *Big Data*, vol. 5, no. 3, pp. 246–255, Sep. 2017.
- [39] G. Volk, J. Gamberding, A. V. Bernuth, and O. Bringmann, "A comprehensive safety metric to evaluate perception in autonomous systems," in *Proc. IEEE 23rd Int. Conf. Intell. Transp. Syst. (ITSC)*, Sep. 2020, pp. 1–8.
- [40] T. Wang, X. Zhu, J. Pang, and D. Lin, "FCOS3D: Fully convolutional one-stage monocular 3D object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops (ICCVW)*, Oct. 2021, pp. 913–922.
- [41] T. Wang, Z. Xinge, J. Pang, and D. Lin, "Probabilistic and geometric depth: Detecting objects in perspective," in *Proc. Conf. Robot Learn.*, 2022, pp. 1475–1485.
- [42] M. Woźniak, M. Wiecek, and J. Siłka, "Deep neural network with transfer learning in remote object detection from drone," in *Proc. 5th Int. ACM Mobicom Workshop Drone Assist. Wireless Commun. 5G Beyond*, Oct. 2022, pp. 121–126.
- [43] Y. Yan, Y. Mao, and B. Li, "SECOND: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, p. 3337, Oct. 2018.
- [44] Z.-Q. Zhao, P. Zheng, S.-T. Xu, and X. Wu, "Object detection with deep learning: A review," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 11, pp. 3212–3232, Nov. 2019.
- [45] X. Zhu, Y. Ma, T. Wang, Y. Xu, J. Shi, and D. Lin, "SSN: Shape signature networks for multi-class object detection from point clouds," in *Proc. Eur. Conf. Comput. Vis.*, A. Vedaldi, H. Bischof, T. Brox, J.-M. Frahm, Eds. Cham, Switzerland: Springer, 2020, pp. 581–597.



**ANDREA CECCARELLI** received the Ph.D. degree in informatics and automation engineering from the University of Florence, Florence, Italy, in 2012. He is currently an Associate Professor of computer science with the University of Florence. His scientific activities originated more than 100 papers that appeared in international conferences, workshops, and journals. His research interests include the design, monitoring, and evaluation of dependable and secure systems, with a preference for experimental approaches. He is a member of the IFIP WG 10.4 on "Dependable Computing and Fault-Tolerance." He has been the Program Committee Co-Chair of the Conference SRDS and the Conference LADC.



**LEONARDO MONTECCHI** (Member, IEEE) received the bachelor's, master's, and Ph.D. degrees in computer science, systems and telecommunications from the University of Florence, Italy, in 2007, 2010, and 2014, respectively. He has been an Associate Professor with the Norwegian University of Science and Technology, Trondheim, Norway, since January 2022. Previously, he was an Assistant Professor with the University of Campinas, Brazil, from 2017 to 2021. His expertise revolves around the modeling of complex systems, including formal models, probabilistic models, and model-driven engineering. His research interests include modeling as a support to the verification and validation of safety-critical and mission-critical systems. He regularly serves as the reviewer for international conferences and journals in the area of systems and software reliability.