
A Step Towards Efficient Evaluation of Complex Perception Tasks in Simulation

Jonathan Sadeghi
Five AI Ltd.

jonathan.sadeghi@five.ai

Blaine Rogers
Five AI Ltd.

James Gunn
Five AI Ltd.

Thomas Saunders
Five AI Ltd.

Sina Samangooei
Five AI Ltd.

Puneet Kumar Dokania
Five AI Ltd.

John Redford
Five AI Ltd.

Abstract

There has been increasing interest in characterising the error behaviour of systems which contain deep learning models before deploying them into any safety-critical scenario. However, characterising such behaviour usually requires large-scale testing of the model that can be extremely computationally expensive for complex real-world tasks. For example, tasks involving compute intensive object detectors as one of their components. In this work, we propose an approach that enables efficient large-scale testing using simplified low-fidelity simulators and without the computational cost of executing expensive deep learning models. Our approach relies on designing an efficient surrogate model corresponding to the compute intensive components of the task under test. We demonstrate the efficacy of our methodology by evaluating the performance of an autonomous driving task in the Carla [10] simulator with reduced computational expense by training efficient surrogate models for PIXOR [44] and CenterPoint [45] LiDAR detectors, whilst demonstrating that the accuracy of the simulation is maintained.

1 Introduction

Deep learning models achieve state of the art performance for a variety of real-world applications [16], however the fact that these models have multiple vulnerabilities such as shift in data distribution [29] and additive perturbations [11, 15, 25, 40], has limited their practical usability in safety-critical situations such as driverless cars. A solution to this problem is to collect and annotate a large diverse dataset that captures all possible scenarios for training and testing. However, since the costs involved in manually annotating such a large quantity of data can be prohibitive, it is often beneficial to employ high-fidelity simulators [17] to potentially produce vast number of diverse scenarios with exact ground-truth annotations at almost no cost.

We would like to use the abundance of annotated samples provided by a simulator to perform *extensive testing* of a given deep learning model. This will allow us to find failure modes of such models before deployment. For example, let us assume that our objective is to find failure modes of a path planner (downstream task g) of a driverless car that takes detected objects as an input from a trained object detector (backbone task f) [7, 16]. Since the failure modes of the detector would have significant impact on the planner, we would like to test the planner by giving inputs directly to the detector [38, 1], and thereby capture dependencies between the detector and the planner. In this setting, one might want to use all possible high-fidelity synthetic generations x from the simulator as an input to the detector to characterise the failure modes of the planner. These failure modes can be identified using metrics for the planner behaviour to highlight data sequences which should be

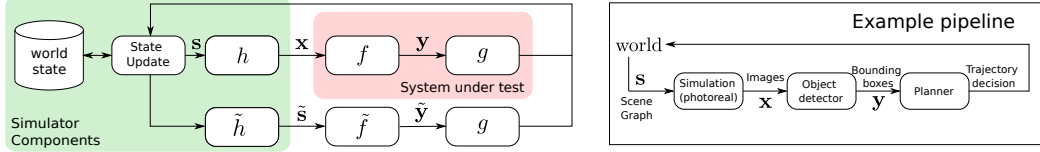


Figure 1: Diagram to show how the simulator, h , and backbone task f can be circumvented during end-to-end testing with a surrogate model \tilde{f} . The outputs of the surrogate model and the backbone task are fed into the model for the downstream task g in the same way because y and \tilde{y} are the elements of the same vector space. The simulator will update the world state s based on the output of the downstream task g , which is used by the expensive simulator h to produce sensor readings x , and efficient low-fidelity simulator \tilde{h} to produce the low-dimensional inputs to the surrogate model \tilde{s} .

inspected by the engineer. However, given that the inference step of an object detector where the input is high-fidelity synthetic data itself is a computationally demanding operation [39], such an approach will not be scalable enough to perform extensive testing of the task.

Motivated by the above example, the goal in this work is to propose an efficient alternative to testing with a high-fidelity simulator and thereby enable large-scale testing. Our approach is to replace the computationally demanding backbone task f with an efficient surrogate \tilde{f} that is trained to mimic the *behaviour* of the backbone task. As opposed to f where the input is a high-fidelity sample x , the input to the surrogate is a much lower-dimensional ‘salient’ variable \tilde{s} . In the example of the object detector and path planner: x might be a high-fidelity simulation of the output of camera or LiDAR sensors, whereas \tilde{s} might be simply the position, orientation and level of occlusion of other vehicles and pedestrians in the scene together with other aspects of the scene like the level of lighting which could also affect the results of the detection function f . The training of the surrogate is performed to provide the following approximate composite model $g(\tilde{f}(\tilde{s})) \approx g(f(x))$. This allows us to efficiently perform rigorous testing on the downstream task using very low-dimensional inputs to an efficient surrogate model for the backbone, as shown in Figure 1. This results in a reduction in compute time whilst maintaining the benefits of using a simulator for testing.

In summary, our contributions are as follows:

- We propose a surrogate modelling approach to enable efficient large-scale testing of complex tasks, and describe one specific methodology for creating such surrogate models in the autonomous driving setting.
- We perform extensive large-scale experiments demonstrating the efficacy of our approach using Carla [10] simulator for driving tasks. We create surrogate models of two well-known LiDAR detectors, PIXOR [44] and Centerpoint [45], as our backbone task.
- We show that our approach is closest to the backbone task compared to the baselines evaluated on several metrics, and yields a 20 times reduction in the compute time.

2 Methodology

Training and testing a model for a complex task primarily involves the following three modules—(1) a **data generator** (h) that provides the high-fidelity sensory input domain \mathcal{X} ; (2) a **backbone-task** (f_θ), parameterised by θ , that maps $x \in \mathcal{X}$ into an intermediate representation $y \in \mathcal{Y}$; and (3) a **downstream task** (g_ϕ), parameterised by ϕ , that takes the intermediate y as an input and maps it into a desired output $z \in \mathcal{Z}$. For example, devising a path planner that takes as input the raw sensory data x from the world via sampler h and outputs an optimal trajectory would rely on intermediate solutions such as accurate detections $f_\theta(x)$ by an object detector f (backbone task) in order to provide the optimal trajectory $g_\phi(f_\theta(x))$. Most real-world problems consist of such complex tasks that heavily depend on intermediate solutions, obtaining which can sometimes be the main bottleneck from both efficiency and accuracy points of view. Considering the same example of a path planner, it is well-known that an object detector is computationally expensive in real-time. Therefore, extensively evaluating the planner that depends on the detector can quickly become computationally infeasible as there exist millions of diverse road scenarios over which the planner should be tested before deployment into any safety-critical environment such as driverless cars in the real-world.

We identify the two following bottlenecks of evaluating extensively such complex tasks: (1) efficiently obtaining a sufficiently diverse set of test scenarios; and (2) efficient inference of the intermediate expensive backbone tasks. Though simulators, as used in this work, can theoretically solve the first problem as they can provide infinitely many test scenarios, their use, in practice, is limited as it still is very expensive for the backbone task f to process these high-fidelity samples obtained from the simulator, and for the simulator to generate these samples.

Our solution to alleviate these bottlenecks is rather simple. Instead of obtaining high-fidelity samples from the simulator h , we generate samples in a low-dimensional vector space which is designed so that these embeddings summarise the crucial information required by the backbone task f to be able to provide accurate predictions. Using these low-dimensional simulator outputs, a rather simplistic and efficient model \tilde{f} can then be trained to mimic the behaviour of the target backbone task f which provides the input for the complex composite task g under test. This will allow very fast and efficient evaluation of the composite task by approximating the inference of the backbone task using a surrogate model. In what follows, we provide details of these approximations.

Obtaining Low-fidelity Simulation Data: The data generation process for simulators is a mapping $h : s \mapsto \mathbf{x}$, where $s \in \mathcal{S}$ denotes the world-state that is normally structured as a scene graph (refer Figure 1). Note, for high-fidelity generations, s is very high dimensional as it contains all the properties of the world necessary to generate a realistic sensor reading \mathbf{x} . For example, in the case of road scenarios, it contains, but is not limited to, positions and types of all vehicles, pedestrians and other moving objects in the world, details of the road shapes and surfaces, surrounding buildings, light and RADAR reflectivity of all surfaces in the simulation, and lighting and weather conditions [10]. There is usually a trade-off between an accurate simulator and the computational expense of the simulator. Therefore, even if a high fidelity simulator is available, it may be intractable to produce sufficient simulated data for training and evaluation as the mapping h itself is expensive.

Noting that a low-fidelity simulator $\tilde{h} : s \mapsto \tilde{s}$ can be created to map high-dimensional s into low-dimensional ‘salient’ variables $\tilde{s} \in \tilde{\mathcal{S}}$ for a variety of tasks [33], in this work, we use \tilde{s} as an input to the surrogate backbone task (design of the surrogate is discussed below). In the simplest case, the mapping $\tilde{h}(\cdot)$ could consist of a subsetting operation. For example, in object detectors, \tilde{h} could output \tilde{s} that contains the position and size of all the actors in the scene. In order to provide more useful information in \tilde{s} , we also include *low-fidelity physical simulations* in \tilde{h} . Specifically, we use a ray tracing algorithm in \tilde{h} to calculate the geometric properties such as occlusion of actors in the scene for one of the ego vehicle’s sensors. Clearly, the subsetting operation and deciding what physical simulations to include in \tilde{h} requires domain knowledge about the backbone task. We believe this is not an unreasonable assumption as for most of the perception related tasks that we are interested in, we have a fairly good idea of what factors are necessary to capture the underlying performance. A more generic setting to automatically learn \tilde{h} could be an interesting future direction.

Efficient Surrogate for the Backbone Task: The next step is to use the low-dimensional \tilde{s} in order to provide reliable inputs for the downstream task. Recall, our objective is to provide an *efficient way to mimic* $f_\theta(h(s))$ so that its output can be passed to the downstream task for large-scale testing (refer Figure 1). To this end, using the low-dimensional \tilde{s} as input, we design and train a surrogate function \tilde{f}_θ such that $\tilde{f}_\theta(\tilde{s}) \approx f_\theta(h(s))$, for all $s \in \mathcal{S}$. By design, the surrogate function takes a very low-dimensional input compared to the high-fidelity \mathbf{x} and, as will be shown in our experiments, is orders of magnitude faster than operating a high-fidelity simulator, h , and the original backbone task, f . The time required for training the surrogate model is typically insignificant compared to the overall simulation budget, and therefore is not discussed in further detail. The performance of the downstream task may differ when the backbone task is replaced with the surrogate model, due to the approximate nature of the surrogate model; erroneous outputs of the surrogate model could be predicted by a full treatment of model uncertainty for the surrogate model, however we do not explore this issue further in the present work.

We now provide the details of the surrogate model. As mentioned, the selection of the salient variables \tilde{s} and the form of the surrogate function is a design choice that requires domain knowledge. Since our experimentation in this work primarily involves testing of a planner that requires an object detector as the backbone task, our choice of salient variables for the input to the detector surrogate includes: *position, linear velocity, angular velocity, actor category, actor size, and occlusion percentage* (we

will specify in each experiment exactly which variables were used). Additional salient variables could also be used, but as a proof of concept of our approach, we chose the above salient variables and did not explore more combinations as these variables already provided promising results. To compute the occlusion percentage, we simulate a low-resolution semantic LiDAR and calculate the proportion of rays terminating in the desired agent’s bounding box [43]. Typically, these salient variables are available at no computational cost when the simulator updates the world-state.

We design a simple probabilistic neural network $\tilde{f}_{\tilde{\theta}}$ as the surrogate for the object detector. To train \tilde{f} , for every \mathbf{s} , a tuple $\{\tilde{\mathbf{s}} = \tilde{h}(\mathbf{s}), \mathbf{y} = f(h(\mathbf{s}))\}$ of input-output is created for every frame, which we process to obtain an input-output tuple for each agent in the scene. Running the perception system to create the training data for the surrogate model incurs an additional computational cost, however we do not consider this to be problematic because the cost is fixed for a particular perception system, and can therefore be amortised over simulations using many different planners. We use the Hungarian algorithm with an intersection over union cost between objects [12] to associate the ground-truth locations and the detections from the *original* backbone task f , on a per-frame basis. Therefore, the training data for the surrogate detector would be $\mathcal{D} = \{\tilde{\mathbf{s}}^i, \tilde{\mathbf{y}}^i\}_{i=1}^k$, and although we have defined \tilde{f} as a function of all ground-truth objects in the scene, in this paper our implementation factorises over each agent in the scene and acts on a single agent basis. This results in a surrogate model which by design cannot predict False Positive detections, as is the case in most of the perception error model literature (see [31, 32, 14, 46, 23]). Choosing a more general surrogate model would avoid this limitation, however, we consider this out of the scope of the present work. Our network architecture for the surrogate is a multi-layered fully-connected network with skip connections, and dropout layers between ‘skip blocks’ (similar to a ResNet [13]), which is shown in the Appendix A. The final layer of the network outputs the parameters of the underlying probability distributions, which are usually a Gaussian distribution (mean and log standard deviation) for the detected position of the objects, and a Bernoulli distribution for the binary valued outputs, e.g. objectness for prediction of false negatives [23]. The training is performed by maximizing the following expected log-likelihood:

$$\mathcal{L}_{\text{total}} = \sum_i \log p(\tilde{\mathbf{y}}_{\text{det}}^i | \tilde{\mathbf{s}}^i) + \mathbb{1}_{\{\mathbf{y}_{\text{det}}^i = 1\}} \log p(\tilde{\mathbf{y}}_{\text{pos}}^i | \tilde{\mathbf{s}}^i), \quad (1)$$

where, associated with the surrogate function $\tilde{f}_{\tilde{\theta}}(\cdot)$, $p(\cdot | \tilde{\mathbf{s}}^i)$ represents the likelihood, $\tilde{\mathbf{y}}_{\text{det}}$ represents the Boolean output which is true if the object was detected, and $\tilde{\mathbf{y}}_{\text{pos}}$ represents a real-valued output describing the centre position of the detected object, respectively. The term $\log p(\tilde{\mathbf{y}}_{\text{det}}^i | \tilde{\mathbf{s}}^i)$ in Eqn. 1 is equivalent to the binary cross-entropy when using a Bernoulli distribution to predict false negatives. Assuming Cartesian components of the positional error to be independent, we obtain:

$$\begin{aligned} \log p(\tilde{\mathbf{y}}_{\text{pos}}^i | \tilde{\mathbf{s}}^i) &= \log \left(\mathcal{N}(\tilde{\mathbf{y}}_{\text{x, pos}}^{(i)}; \mu_x, \sigma_x) \cdot \mathcal{N}(\tilde{\mathbf{y}}_{\text{y, pos}}^{(i)}; \mu_y, \sigma_y) \right) \\ &= \frac{(\tilde{\mathbf{y}}_{\text{x, pos}}^{(i)} - \mu_x)^2}{2\sigma_x^2} + \frac{(\tilde{\mathbf{y}}_{\text{y, pos}}^{(i)} - \mu_y)^2}{2\sigma_y^2} + \log \sigma_x \sigma_y \end{aligned} \quad (2)$$

where μ and $\log(\sigma)$ are the outputs of the fully connected neural network. For further details on how probabilistic neural networks can be trained for multiple target variables, see [20, 21].

3 Related work

End-to-end evaluation in simulation: End-to-end evaluation refers to the concept of evaluating components of a modular machine learning pipeline together in order to understand the performance of the system as a whole. Such approaches often focus on strategies to obtain equivalent performance using a lower fidelity simulator whilst maintaining accuracy to make the simulation more scalable [33, 3, 35]. Similarly, Wang et al. [42] use a realistic LiDAR simulator to modify real-world LiDAR data which can then be used to search for adversarial traffic scenarios to test end-to-end autonomous driving systems. Kadian et al. [19] attempt to validate a simulation environment by showing that an end-to-end point navigation network behaves similarly in the simulation to in the real-world by using the correlation coefficient of several metrics in the real-world and the simulation. End-to-end testing is possible without a simulator, for example Philion et al. [30] evaluate the difference between planned vehicle trajectories when planning using ground truth and a perception system and show that this enables important failure modes of the perception system to be identified. End to end testing could be

performed more efficiently using knowledge distillation or network quantisation [8]. However, such approaches would still require running h to produce \mathbf{x} , which is typically computationally expensive and therefore we do not perform further comparison with these approaches.

Our approach differs from these in that the surrogate model methodology enables end-to-end evaluation without running the backbone model in the simulation.

Perception error models (PEMs): Perception Error Models (PEMs) are surrogate models used in simulation to replicate the outputs of perception systems so that downstream tasks can be evaluated as realistically as possible. Piazzoni et al. [31] present a PEM for the pose and class of dynamic objects, where the error distribution is conditioned on the weather variables, and use the model to validate an autonomous vehicle system in simulation on urban driving tasks. Piazzoni et al. [32] describe a similar approach using a time dependent model and a model for false negative detections. Time dependent perception PEMs have also been used by Berkhahn et al. [5] to model traffic intersections with a behaviour model and a stochastic process misperception model on velocity, and Hirsenkorn et al. [14] by creating a Kernel Density Estimator model of a filtered radar sensor, where the simulated sensor is modelled by a Markov process. Zec et al. [46] propose to model an off the shelf perception system using a Hidden Markov Model. Modern machine learning techniques have also been used to create PEMs, for example Krajewski et al. [23] create a probabilistic neural network model for a LiDAR sensor, Arnelid et al. [2] use Recurrent Conditional Generative Adversarial Networks to model the output of a fused camera and radar sensor system, and Suhre and Malik [37] describe an approach for simulating a radar sensor using conditional variational auto-encoders.

Perception Error Models should not be confused with surrogate models which approximate the entire system simulation, i.e. $g(f(h(\mathbf{s})))$ over all time steps, rather than the perception system specifically. These full-system surrogate models are used in simulation to efficiently search for failures in autonomous driving systems [4, 36, 9, 41].

Our model is most similar to Mitra et al. [26], where an autoregressive neural network PEM is evaluated in simulation by visually comparing ego behaviour to the behaviour under no perception errors. We describe a more general framework for the training of surrogate models in a probabilistic machine learning context with a larger-scale evaluation than the other papers in this section.

4 Experiments

Overview: We use Carla simulator [10] to analyze the behaviour of an agent in two driving tasks: (1) adaptive cruise control (ACC) and (2) the Carla leaderboard. The Carla configuration is provided in Appendix C. The agent uses a LiDAR object detector f (backbone task) to detect other agents and make plans accordingly. Using the methodology described in Section 2, we construct a **Neural Surrogate (NS)** \tilde{f} that, as opposed to f , does not depend on high-fidelity simulated LiDAR scans. We show that the surrogate agents behave similarly to the real agent while being extremely efficient.

Details of Driving Scenarios:

1. *Adaptive Cruise Control (ACC):* The agent follows a fast moving vehicle. Suddenly, the fast moving vehicle cuts out into an adjacent lane, revealing a parked car. The agent must brake to avoid a collision. See [18] for a review of ACC.
2. *Carla Leaderboard:* See Ros et al. [34] for details. This contains a far more diverse set of driving scenarios, including urban and highway driving, and is approximately a two orders of magnitude increase in the total driving time relative to the ACC task. Therefore, this evaluation can be seen as a large scale evaluation of our methodology.

For the ACC task we use a very simple planner described in Section 4.1 that maintains a constant speed and brakes to avoid obstacles. For the more demanding Carla Leaderboard we use a more robust planner and detector, described in detail in Section 4.2.

Baselines: We compare our approach Neural Surrogate (NS) against three strong baseline surrogate models (\tilde{f}), which are similar to those used in the PEM literature described in Section 3:

- Planning using *ground truth (GT)* locations of agents, which are available from $\tilde{\mathbf{s}}$.

- A *logistic regression (LR)* surrogate is trained to predict the false negatives (missed detections) of the backbone task f for a given \tilde{s} . Note, \tilde{s} here is same as the salient variables used in NS. The logits for the true class probability are then predicted by first passing the input via a linear mapping W and then applying a sigmoid function, i.e. $p_t = \sigma(W\tilde{s})$ where σ is the sigmoid function. The model is trained using the focal loss, $\mathcal{L}(p_t) = -\alpha_t(1 - p_t)^\gamma \log(p_t)$ where p_t represents the true class probability and $\alpha_t = \mathbb{1}_{\{y_{\text{det}}^i=1\}}\alpha + \mathbb{1}_{\{y_{\text{det}}^i=0\}}(1 - \alpha)$, in order to account for class imbalance [27, 24]. The hyperparameters used for the focal loss were manually tuned to be $\alpha = 0.6$ and $\gamma = 2$ via cross-validation over the classification metrics. Once a box has been identified as a true-positive detection of f by the LR model, its **exact** co-ordinates from \tilde{s} are then passed to the planner.
- A *Gaussian Fuzzer (GF)* is a simple surrogate model where the exact position and velocity of a box obtained from \tilde{s} are simply **perturbed** with samples from independent Gaussian and StudentT distributions respectively (StudentT distributions are used due to the heavy tails of the detected velocity errors). This is Eqn. 2 with fixed μ and σ , i.e. not a function of other variables in \tilde{s} . These parameters are obtained analytically using Maximum Likelihood Estimation (MLE). For example, for the Gaussian distribution over positional errors, the MLE solution is simply the empirical mean and the standard deviation of the detection position errors which is obtained using the train set [6].

In the Carla leaderboard evaluation, only the ground truth baseline is used. The hyperparameters used for the training of all the surrogate models are shown in Appendix B.

Surrogate Training Data: In both experiments, the Carla leaderboard scenarios are used to obtain training data for surrogate models and the LiDAR detector. The dataset consists of driving scenarios which are approximately 10 minutes long, featuring a variety of vehicles, cyclists and pedestrians in a mixture of urban and highway scenes.

1. In the first experiment, scenarios 0-9 are used for training and scenario 10 is used for testing. Pedestrians are excluded from the data because the ACC task does not involve pedestrians.
2. The second experiment contains a wider variety of driving scenarios so the collected dataset is larger; scenarios 0-58 are used for training and scenarios 59-62 are used for testing.

Metrics: We use common classification and regression metrics to directly compare the outputs of the surrogate model and the real model on the backbone task:

- Precision and recall.
- Classification accuracy.
- Sampled position mean squared error (spMSE): the mean squared error of the detections y or surrogate predictions \tilde{y} , as appropriate, relative to the ground truth values in s .

We wish to quantify (1) how closely the surrogate mimics the backbone task f ; and (2) how close it is to the ground-truth obtained from \tilde{s} . While evaluating a surrogate model *relative to* f , a false negative of \tilde{f} would be a situation when an agent is detected by \tilde{f} which was in fact missed; conversely evaluating a surrogate model *relative to the ground truth* means that a false negative of \tilde{f} would be when an agent is not detected by \tilde{f} which is in fact present in the ground truth data. When we evaluate surrogate models *relative to the detector* (comparing y to \tilde{y}), the best surrogate is the one with the highest value of the evaluation metric. However, when we evaluate surrogate models *relative to the ground truth* (comparing y or \tilde{y} , as appropriate, to s), the best surrogate is the one whose score is closest to the detector’s score. The metrics are only evaluated for objects within 50m of the ego vehicle, since objects further than this are unlikely to influence the ego vehicle’s behaviour.

We use other metrics to compare the performance of the surrogate and real agents on the downstream task. For the ACC task we evaluate the runtime per frame (with and without h or \tilde{h}), Maximum Braking Amplitude (MBA), and MBA timestamp (tMBA) relative to the start of the scenario. MBA quantifies the degree to which the braking was applied relative to the maximum possible braking.

We also evaluate the mean Euclidean norm (meanEucl), which we define as the time integrated norm of the stated quantity, i.e. to compare variables $v_1(t)$ and $v_2(t)$, the metric is

$$\frac{1}{t'} \int_{t=0}^{t=t'} \|v_1(t) - v_2(t)\|_2 dt. \quad (3)$$

Table 1: Metrics comparing surrogates \tilde{f} (LR, NS, and GF), and the PIXOR detector (the backbone task f) on the train and test set *relative to the ground-truth*. The surrogate closest to the detector is shown in bold. Note, surrogates trivially achieve a precision of 1 as they do not model false positives.

	NS (our)	GF	LR	Detector	NS (our)	GF	LR	Detector
	Train Set				Test Set			
Precision	1	1	1	0.934	1	1	1	0.842
Recall	0.642	1	0.218	0.614	0.533	1	0.238	0.476
spMSE	0.231	0.294	0	0.239	0.267	0.295	0	0.273

Table 2: Metrics comparing LR, NS (our) and GF surrogate models on the train and test set *relative to the detector*. The GF trivially achieves a recall of 1, since the detector does not model false negatives.

	LR	GF	NS (our)	LR	GF	NS (our)
	Train Set			Test Set		
Accuracy \uparrow	0.688	0.244	0.954	0.784	0.176	0.915
Recall \uparrow	0.277	1	0.931	0.412	1	0.823
Precision \uparrow	0.332	0.244	0.886	0.392	0.176	0.730

We evaluate Eqn. 3 with a discretised sum. This metric is a natural, time dependent, method of *comparing trajectories* in Euclidean space. In Appendix G, we provide a relationship between Eqn. 3, and the planner KL-divergence metric proposed by Phillion et al. [30]. We also compute the maximum Euclidean norm (maxEucl) to show the maximum instantaneous difference in the stated quantity,

$$\max_{t \in [0, t']} ||v_1(t) - v_2(t)||_2. \quad (4)$$

In the Carla leaderboard task we compare the metrics used for Carla leaderboard evaluation i.e. *route completion, pedestrian collisions and vehicle collisions* for the detector and NS. We also compute the cumulative distribution functions of the time between collisions for the detector, NS, and GT.

4.1 Carla Adaptive Cruise Control Task

Configuration Details: In this experiment, our backbone task f consists of a PIXOR LiDAR detector [44] trained on simulated LiDAR pointclouds from Carla [44], followed by a Kalman filter which enables the calculation of the velocity of objects detected by the LiDAR detector. Therefore, \mathbf{y} and $\tilde{\mathbf{y}}$ consist of position, velocity, agent size and a binary variable representing detection of the object. To simplify the surrogate model, in this particular experiment, we assume that the ground-truth value of the agent size can be used by the planner whenever required. The salient variables $\tilde{\mathbf{s}}$ consist of position, orientation, velocity, angular velocity, object extent, and percentage occlusion. The downstream task consists of a planner which is shown in further detail in Appendix D.1.

Results: Regression and classification performance *relative to the ground-truth* on the train and test set are shown in Table 1 for both the surrogate models and the detector. Table 2 shows similar metrics to Table 1, but this time computed for the surrogate models *relative to the detector*. This shows that although the LR surrogate is predicting a similar proportion of missed detections, the NS is more effective at predicting these when the detector would also have missed the detection. Appendix E contains plotted empirical cumulative distribution functions for the positional and velocity error predicted by each surrogate model and the true detector relative to the true object locations.

MBA and time efficiency results are shown in Table 3. The surrogates are multiple times faster than the backbone task while showing MBA behaviour similar to the backbone task. Notably the wall-time taken per step (DTPF) is about **100 times higher** for the PIXOR Detector than the surrogate models, excluding the simulator rendering time, with all models running on an Intel Core i7-8750H CPU. Including the simulator rendering time, the difference is reduced to **20 times** (TTPF), indicating that the majority of the time savings are achieved by removing the object detector from the simulation pipeline. The TTPF for GF is approximately 0.06 seconds less than for the other surrogate models since in this case the headless simulator \tilde{h} does not have to calculate the occlusion of agents. In real-world applications, detectors typically run on GPU hardware which results in much quicker

Table 3: Timestamp of Maximum Braking Amplitude (tMBA), MBA, Detector time / frame for f and \tilde{f} (excluding simulator rendering time) (DTPF), and Total Time / Frame (TTPF) – including simulator time (h, \tilde{h}), state update, and the planner time (g). For all metrics the most similar model to the detector is shown in bold. All time noted in seconds.

	Det.	GF	NS (our)	LR	GT
tMBA	13.5	19.0	13.5	15.8	19.9
MBA	1	0.378	1	0.518	0.093
DTPF	1.511	0.013	0.017	0.013	0.005
TTPF	1.644	0.023	0.087	0.083	0.022

Table 4: Mean and Max Euclidean norm of displacement, computed for the ego position and velocity traces separately (Eqn. 3 and Eqn. 4). For both metrics the most similar model to the detector f is shown in bold. The normalised value of the metric, obtained by dividing the metric by the (GT, Det). value, is shown in brackets.

	Det.	GF	NS (our)	LR	GT		Det.	GF	NS (our)	LR	GT
	meanEucl Position ↓						maxEucl Position ↓				
Det.	0 (0.0)	4.9 (0.98)	0.40 (0.08)	1.6 (0.32)	4.9 (1.0)		0.0 (0.0)	19 (0.99)	2.2 (0.11)	7.9 (0.41)	19 (1.0)
GF		0 (0.0)	4.5 (0.92)	4.1 (0.83)	0.15 (0.03)			0.0 (0.0)	17. (0.87)	11. (0.57)	0.36 (0.019)
NS			0 (0.0)	1.2 (0.24)	4.6 (0.93)				0.0 (0.0)	5.7 (0.30)	17. (0.89)
LR				0 (0.0)	4.2 (0.85)					0.0 (0.0)	11. (0.59)
GT					0 (0.0)						0.0 (0.0)
	meanEucl Velocity ↓						maxEucl Velocity ↓				
Det.	0 (0.0)	1.5 (0.99)	0.20 (0.13)	0.71 (0.46)	1.5 (1.0)		0.0 (0.0)	5.5 (0.99)	2.0 (0.35)	4.1 (0.73)	5.5 (1.0)
GF		0 (0.0)	1.3 (0.87)	0.93 (0.60)	0.027 (0.017)			0.0 (0.0)	4.7 (0.84)	3.4 (0.61)	0.11 (0.019)
NS			0 (0.0)	0.52 (0.34)	1.3 (0.88)				0.0 (0.0)	3.2 (0.57)	4.7 (0.85)
LR				0 (0.0)	0.94 (0.61)					0.0 (0.0)	3.4 (0.62)
GT					0 (0.0)						0.0 (0.0)

inference. However, for simulation purposes, removing the requirement of using a GPU for perception inference has the potential to significantly reduce costs.

In Table 4, a selection of pairwise metrics is shown comparing the ego trajectory in each simulation environment. The pairwise metrics show that using a surrogate model produces closer agent behaviour to the backbone task (LiDAR detector) compared to GT, both for metrics based on velocity and position. The NS is the best performing model on all pairwise metrics. Plots of the actors’ trajectories are shown in Appendix F. The GF produces similar ego trajectories to the GT baseline, and this is most likely because false negatives, which cause delayed braking and are therefore influential in this scenario, are not included in both cases. The metrics indicate that the LR model is most similar to the NS, however, the ego trajectories produced by the LR are less similar to those produced by the LiDAR detector than those produced by the NS.

4.2 Carla Leaderboard Evaluation

Configuration Details: In this experiment, the backbone task f is a CenterPoint LiDAR detector [45] for both vehicles and pedestrians, trained on the simulated data from Carla in addition to proprietary real-world data. The downstream planner g is a modified version of the BasicAgent from the Carla Python API, which follows a fixed path and performs emergency braking to avoid vehicle and pedestrian collisions (described in further detail in Appendix D.2).

The NS model architecture is mostly the same as in Section 2, but the agent velocity is removed from y , since the BasicAgent does not require the velocities of other agents. In addition, an extra salient variable is provided to the network in \tilde{s} : a one hot encoding of the class of the ground truth object (vehicle or pedestrian) and in the case of the object being a vehicle, the make and model of the vehicle. Since the training dataset is imbalanced and contains more vehicles at large distances from the ego vehicle, minibatches for the training are created by using a stratified sampling strategy: the datapoints are weighted using the inverse frequency in a histogram over distance with 10 bins, resulting in a balanced distribution of vehicles over distances.

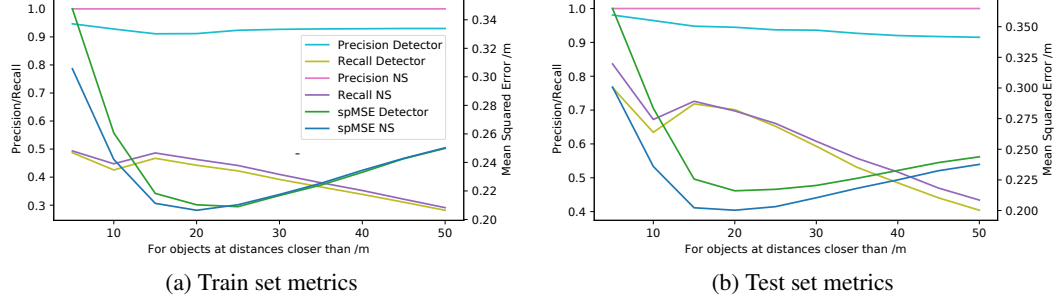


Figure 2: Metrics comparing NS and the CenterPoint detector on train and test set *relative to the ground-truth*. The precision of NS is always 1 because, by design, it is unable to predict false positives (we do not re-associate the detections to ground truth after they are generated by the NS, and instead perform the evaluation using the association used to generate the boxes).

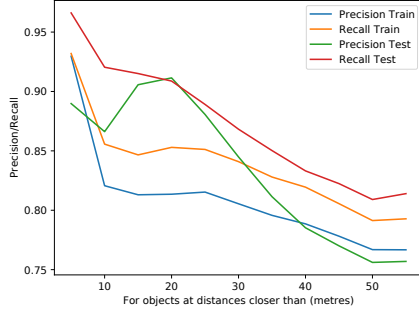


Figure 3: Comparing NS *relative to the detector*.

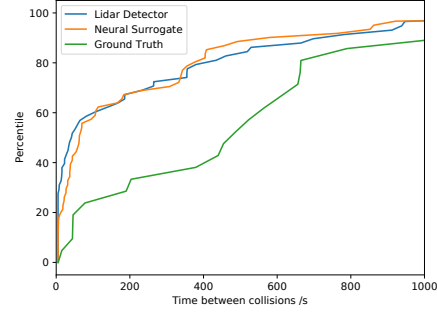


Figure 4: Cumulative distribution of the time between collisions for Carla leaderboard evaluation for NS, GT, and the CenterPoint.

Results: Metrics on the train and test set *relative to the ground-truth* are shown in Figure 2 for both the NS and the detector. The NS closely mimics the behaviour of the CenterPoint detector with respect to these metrics. Figure 3 shows similar metrics computed for the NS *relative to the detector*. Again, the NS performs well on all metrics, particularly for objects at lower distances which are usually the most influential on the behaviour of ego.

Metrics used for Carla leaderboard evaluation are summarised in Table 5. Since the NS does not model false positive detections, the route completion is lower in some scenarios where a false positive LiDAR detection of street furniture confuses the planner, which does not happen for the NS or the GT. We leave the surrogate modelling of false positive detections as a future research direction. Figure 4 shows cumulative distribution functions of the time between collisions; NS is *clearly more similar* to the LiDAR detector than the ground truth. The NS captures the time between collisions for the CenterPoint model much more effectively than the GT, and at a fraction of the cost of running the high-fidelity simulator and the CenterPoint detector. Appendix H shows examples of LiDAR detector errors reproduced by the surrogate model.

Table 5: Carla leaderboard evaluation metrics.

	LiDAR Det. (CenterPoint)	NS (our)	GT
# Ped. Collisions	1	1	1
# Veh. Collisions	60	63	23
Avg Completion (%)	81.870	93.933	91.373
Median time between collisions (seconds)	41.250	62.800	470.950

5 Conclusions

In this paper, we proposed an approach to enable efficient large-scale testing of complex perceptual tasks in a simulated environment. We showed that it is possible to create an efficient surrogate model corresponding to heavy-compute components (for example, the backbone task of detecting objects) of a complex task such that the input now is much lower-dimensional, and the inference is multiple times faster. We provided extensive analysis to show that such surrogate models, while showing similar behaviour to their heavy-compute counterparts when compared using variety of metrics (precision, recall, trajectory similarity, etc.), were multiple times faster as well.

In future work, more complex surrogate models could be designed which behave more similarly to the baseline model, and the effect of surrogate model uncertainty and calibration studied. In addition modelling of false positive detections could be added to the surrogate model. This is a more complex problem than false negative detections as additional non-obvious context information will need to be added to the salient variables, because false positives are often caused by factors external to the agents present in the scene.

Acknowledgments and Disclosure of Funding

We are grateful to all colleagues at Five who have contributed to insightful discussions about Perception Error Models, including Philip Torr, Andrew Blake, Simon Walker, John McDermid, Sebastian Kaltwang, Adam Charytoniuk, Torran Elson and Romain Mueller. We are particularly grateful to Romain Mueller, Luca Bertinetto, Francesco Pinto and Anuj Sharma for proofreading this manuscript.

References

- [1] Afsoon Afzal, Deborah S Katz, Claire Le Goues, and Christopher S Timperley. A study on the challenges of using robotics simulators for testing. *arXiv preprint arXiv:2004.07368*, 2020.
- [2] Henrik Arnelid, Edvin Listo Zec, and Nasser Mohammadiha. Recurrent conditional generative adversarial networks for autonomous driving sensor modelling. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 1613–1618. IEEE, 2019.
- [3] Aravind Balakrishnan. Closing the modelling gap: Transfer learning from a low-fidelity simulator for autonomous driving. Master’s thesis, University of Waterloo, 2020.
- [4] Halil Beglerovic, Michael Stolz, and Martin Horn. Testing of autonomous vehicles using surrogate models and stochastic optimization. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6, 2017. doi: 10.1109/ITSC.2017.8317768.
- [5] Volker Berkhahn, Marcel Kleiber, Johannes Langner, Chris Timmermann, and Stefan Weber. Traffic dynamics at intersections subject to random misperception. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–11, 2021. doi: 10.1109/TITS.2020.3045480.
- [6] C.M. Bishop. *Pattern Recognition and Machine Learning*. Information science and statistics. Springer-Verlag New York, 2006. ISBN 978-0-387-31073-2.
- [7] NTS Board. Collision between vehicle controlled by developmental automated driving system and pedestrian. *Nat. Transpot. Saf. Board, Washington, DC, USA, Tech. Rep. HAR-19-03*, 2019. URL <https://www.nts.gov/investigations/AccidentReports/Reports/HAR1903.pdf>.
- [8] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*, 2017.
- [9] Anthony Corso, Robert Moss, Mark Koren, Ritchie Lee, and Mykel Kochenderfer. A survey of algorithms for black-box safety validation of cyber-physical systems. *J. Artif. Int. Res.*, 72:377–428, October 2021. ISSN 1076-9757. doi: 10.1613/jair.1.12716.
- [10] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16, 2017.

- [11] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1625–1634, 2018.
- [12] David A Forsyth and Jean Ponce. *Computer vision: a modern approach*. Pearson, 2012.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Nils Hirsenkorn, Timo Hanke, Andreas Rauch, Bernhard Dehlink, Ralph Raschhofer, and Erwin Biebl. Virtual sensor models for real-time applications. *Advances in Radio Science*, 14:31–37, 2016.
- [15] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 125–136, 2019.
- [16] Joel Janai, Fatma Güney, Aseem Behl, Andreas Geiger, et al. Computer vision for autonomous vehicles: Problems, datasets and state of the art. *Foundations and Trends® in Computer Graphics and Vision*, 12 (1–3):1–308, 2020.
- [17] Matthew Johnson-Roberson, Charles Barto, Rounak Mehta, Sharath Nittur Sridhar, Karl Rosaen, and Ram Vasudevan. Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 746–753. IEEE, 2017.
- [18] Ronald K Jurgen. Adaptive cruise control. Technical report, SAE Technical Paper, 2006.
- [19] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. Are we making real progress in simulated environments? measuring the sim2real gap in embodied visual navigation. *arXiv preprint arXiv:1912.06321*, 2019.
- [20] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584, 2017.
- [21] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] Robert Krajewski, Michael Hoss, Adrian Meister, Fabian Thomsen, Julian Bock, and Lutz Eckstein. Neural-networks-based modeling of automotive perception errors using drones as reference sensors.
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [25] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- [26] Pallavi Mitra, Apratirn Choudhury, Vimal Rau Aparow, Giridharan Kulandaivelu, and Justin Dauwels. Towards modeling of perception errors in autonomous vehicles. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3024–3029. IEEE, 2018.
- [27] Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip Torr, and Puneet Dokania. Calibrating deep neural networks using focal loss. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 15288–15299, 2020.
- [28] Ouster, Inc. OS1Mid-Range High-Resolution Imaging Lidar. URL <http://data.ouster.io/downloads/OS1-lidar-sensor-datasheet.pdf>.
- [29] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in Neural Information Processing Systems*, 32:13991–14002, 2019.

- [30] Jonah Philion, Amlan Kar, and Sanja Fidler. Learning to evaluate perception models using planner-centric metrics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14055–14064, 2020.
- [31] Andrea Piazzoni, Jim Cherian, Martin Slavik, and Justin Dauwels. Modeling sensing and perception errors towards robust decision making in autonomous vehicles. *arXiv preprint arXiv:2001.11695*, 2020.
- [32] Andrea Piazzoni, Jim Cherian, Martin Slavik, and Justin Dauwels. Modeling perception errors towards robust decision making in autonomous vehicles. In *IJCAI*, 2020.
- [33] Samira Pouyanfar, Muneeb Saleem, Nikhil George, and Shu-Ching Chen. Roads: Randomization for obstacle avoidance and driving in simulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [34] German Ros, Vladfen Koltun, Felipe Codevilla, and Antonio Lopez. The CARLA Autonomous Driving Challenge. <https://carlachallenge.org/>, 2019.
- [35] Ahmad El Sallab, Ibrahim Sobh, Mohamed Zahran, and Nader Essam. Lidar sensor modeling and data augmentation with gans for autonomous driving. *arXiv preprint arXiv:1905.07290*, 2019.
- [36] Aman Sinha, Matthew O’Kelly, Russ Tedrake, and John C Duchi. Neural bridge sampling for evaluating safety-critical autonomous systems. *Advances in Neural Information Processing Systems*, 33, 2020.
- [37] Alexander Suhre and Waqas Malik. Simulating object lists using neural networks in automotive radar. In *2018 19th International Conference on Thermal, Mechanical and Multi-Physics Simulation and Experiments in Microelectronics and Microsystems (EuroSimE)*, pages 1–5. IEEE, 2018.
- [38] Ardi Tampuu, Tarmet Matiisen, Maksym Semikin, Dmytro Fishman, and Naveed Muhammad. A survey of end-to-end driving: Architectures and training methods. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [39] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10781–10790, 2020.
- [40] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, number 2019, 2019.
- [41] Jonathan Uesato, Ananya Kumar, Csaba Szepesvari, Tom Erez, Avraham Ruderman, Keith Anderson, Krishnamurthy Dj Dvijotham, Nicolas Heess, and Pushmeet Kohli. Rigorous agent evaluation: An adversarial approach to uncover catastrophic failures. In *International Conference on Learning Representations*, 2018.
- [42] Jingkang Wang, Ava Pun, James Tu, Sivabalan Manivasagam, Abbas Sadat, Sergio Casas, Mengye Ren, and Raquel Urtasun. Advsim: Generating safety-critical scenarios for self-driving vehicles. *CoRR*, abs/2101.06549, 2021.
- [43] Yu Xiang and Silvio Savarese. Object detection by 3d aspectlets and occlusion reasoning. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 530–537, 2013.
- [44] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018.
- [45] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. *arXiv:2006.11275*, 2020.
- [46] Edwin Listo Zec, Nasser Mohammadiha, and Alexander Schliep. Statistical sensor modelling for autonomous driving using autoregressive input-output hmms. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 1331–1336. IEEE, 2018.

A Network Architecture

Our network architecture for the surrogate is a multi-layered fully-connected network with skip connections, and dropout layers between ‘skip blocks’ (similar to a ResNet [13]), which is shown in the Figure 5. The final layer of the network outputs the parameters of the underlying probability distributions.

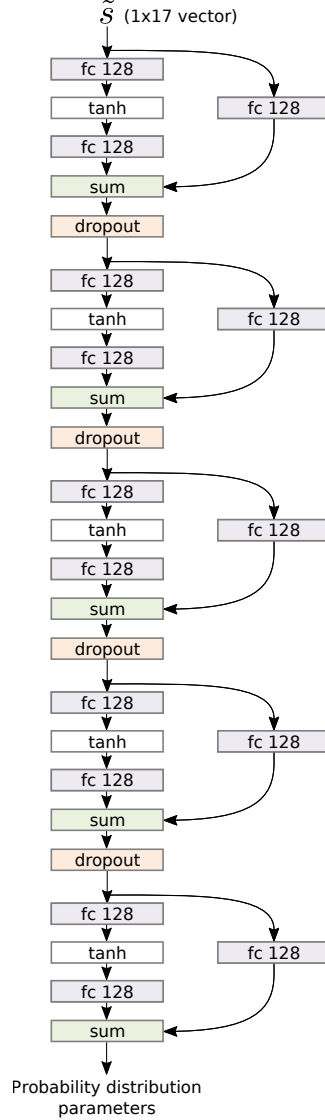


Figure 5: Diagram of architecture used for Neural Surrogate. Fully connected layers are shown as fc blocks where the number represents the output channel size.

B Experimental Hyperparameters

The hyperparameters used to train the surrogate models in the ACC experiment are shown in Table 6. Adam optimiser was used for all training [22]. Hyperparameters not shown were set to default values. The hyperparameters were selected by manual tuning.

The hyperparameters used to train the neural network for the carla leaderboard evaluation are shown in Table 7.

Table 6: Hyperparameters used to train the surrogate models in ACC experiment.

	Neural Surrogate	Logistic Regression	Gaussian Fuzzer
Learning Rate	1e−3	1e−2	1e−2
$N_{\text{iterations}}$	20000	3000	3000
Dropout Rate	0.3	N/A	N/A
Batch Size	128×128	16×128	16×128

Table 7: Hyperparameters used to train the surrogate models in Carla leaderboard evaluation.

	Neural Surrogate
Learning Rate	1e−4
$N_{\text{iterations}}$	50000
Dropout Rate	0.0
Batch Size	12000
γ (Focal Loss)	2.0
α (Focal Loss)	0.5

C Carla Simulator Configuration

The settings for the Carla lidar sensor are shown in Table 8.

Table 8: Lidar sensor configuration. The configuration is set to be approximately equal to an Ouster lidar sensor [28].

Property	Value
Sensor Position	$x = 0.9, y = 0.0, z = 1.8$
Sensor Orientation	roll = 0.0, pitch = 0.0, yaw = 0.0
Sensor Type	<code>sensor.lidar.ray_cast</code>
Rotation Frequency	20
Channels	128
Points per second	2 620 000
Upper Field of View	11.25
Lower Field of View	-11.25
Range	100
Atmosphere attenuation rate	0.004
Noise stddev	0.1
Dropoff general rate	0.45
Dropoff intensity limit	0.8
Dropoff zero intensity	0.9

D Planners

This appendix describes the ACC and BasicAgent planners used in the experiments in Section 4.1 and Section 4.2 respectively.

D.1 ACC Planner Pseudo Code

A PID controller is used in combination with the planner in Listing 1 to control the vehicle throttle and brake. The planner accelerates ego to a maximum velocity unless a slow moving vehicle is detected in the same lane as ego, in which case ego will attempt to decelerate so that ego’s velocity matches that of the slow moving vehicle. If the ego is closer than 0.1 metres to the slow moving vehicle, then it applies emergency braking.

```

slow_threshold = 5
safe_stop_headway = 15
forward_horizon = 100
cruise_speed = 13.9
last_target_speed = cruise_speed
lane_width = 4.5
time_threshold = 0.5
timestep = 0.05

in_lane = filter(lambda x: abs(x.position.y) < lane_width / 2, objects)
within_horizon = filter(lambda x: x.position.x < forward_horizon, in_lane)
slow = list(filter(lambda x: x.speed < slow_threshold, within_horizon))

if len(slow) > 0:
    closest_agent = sorted(slow, key=lambda x: x.position.x)[0]

    distance_to_closest = distance_to_closest_agent_vertex(closest_agent)
    distance_to_safe = distance_to_closest - safe_stop_headway

    if distance_to_safe < 0.1:
        target_speed = 0
    else:
        speed_of_closest = closest_agent.speed

        if speed_of_closest < 0.5:
            speed_of_closest = 0

        accel = (speed_of_closest ** 2 - last_target_speed ** 2) / (2 * distance_to_safe)
        target_speed = last_target_speed + accel * timestep

else:
    target_speed = last_target_speed + target_accel * timestep

```

Listing 1: Planner Pseudo Code

D.2 BasicAgent Modifications

The BasicAgent planner uses a PID controller to accelerate the vehicle to a maximum speed, and stops only if a vehicle within a semicircle of specific radius in front of ego is detected where the vehicle’s centre is in the same lane as ego. We made changes to the BasicAgent planner to improve its performance. We modified BasicAgent to avoid pedestrian collisions and brake when the corner of a vehicle is inside a rectangle of lane width in front of the ego such that the vehicle’s lane is the same as one of ego’s future lanes. Also, the BasicAgent was modified to drive slower close to junctions.

E Surrogate model error distributions

Empirical cumulative distribution functions for the positional and velocity error predicted by each surrogate model and the true detector relative to the true object locations are shown in Figure 6.

F Ego Behaviour Diagnostics

The diagnostics for the ACC experiment are shown in Figures 7, 8, 9, 10 and 11.

G Comparison with Planner KL Divergence

In this appendix, we explain the relationship between the mean Euclidean norm metric, and the Planner KL divergence metric proposed by Phillion et al. [30].

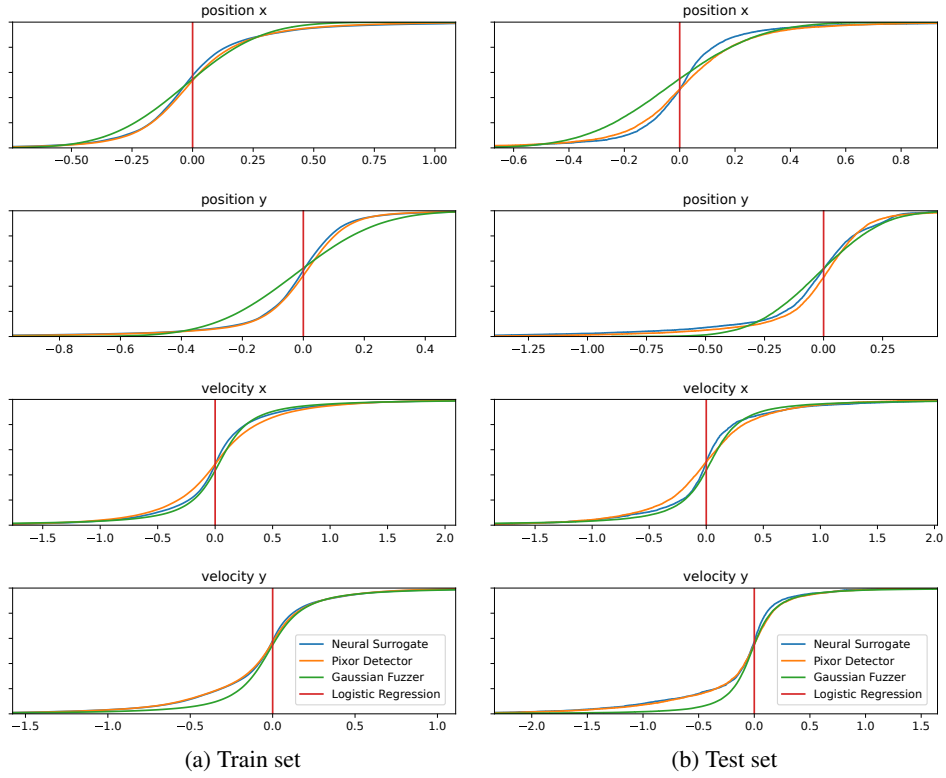


Figure 6: Surrogate model fitted error distributions on train and test set. Note that in all cases the Neural Surrogate (our) distribution is closer than any of the other baselines to the distribution of the PIXOR Detector being approximated. Although the fit of the models to the training set appears similar from these plots, the performance of a planner downstream from these models is both visually and quantifiably different.

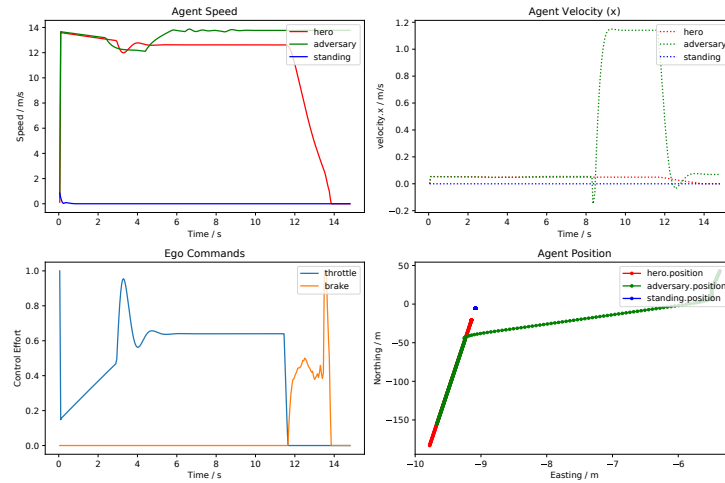


Figure 7: Diagnostics for simulation with the full upstream detector task.

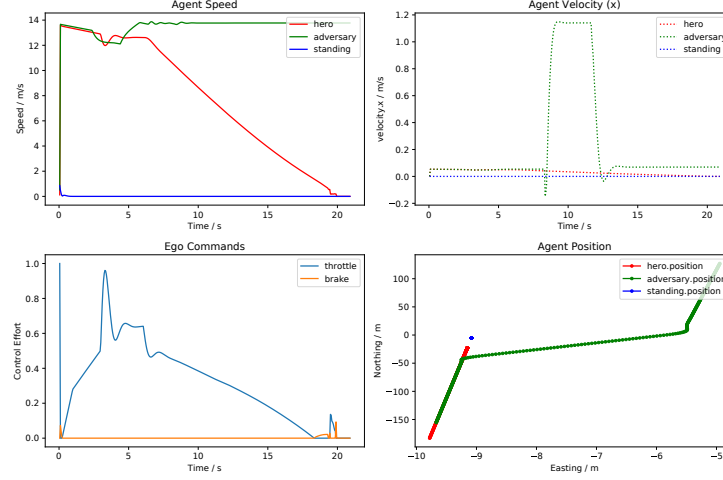


Figure 8: Diagnostics for simulation with the upstream detector outputs substituted for ground truth values.

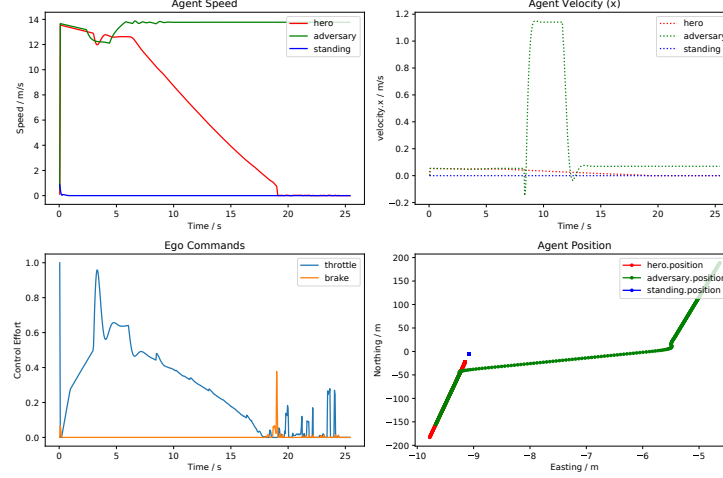


Figure 9: Diagnostics for simulation with the upstream detector outputs substituted for values generated by a simple 'Fuzzer' surrogate model.

The KL divergence between the plan produced by an agent planning based on a detector and a surrogate model is given by:

$$\text{P-KL} = \mathbb{E}_{p(\mathbf{z}^1|f(\mathbf{x}))} \left(\log \frac{\prod_{t=1}^{t=T} p(\mathbf{z}_t^1|f(\mathbf{x}))}{\prod_{t=1}^{t=T} \int p(\mathbf{z}_t^1|\mathbf{y})p(\mathbf{y}|\tilde{\mathbf{s}})d\mathbf{y}} \right), \quad (5)$$

where \mathbf{z}_t^1 is the position of ego at timestamp t , $\mathbf{z}^1 = \{\mathbf{z}_1^1, \dots, \mathbf{z}_T^1\}$, $f : \mathbf{x} \mapsto \mathbf{y}$ is the detector, $p(\mathbf{z}_t^1|\mathbf{y})$ is the probabilistic planner, and $p(\mathbf{y}|\tilde{\mathbf{s}})$ is the probability distribution associated with the surrogate model for the detector, i.e. $\tilde{f}(\mathbf{y}|\tilde{\mathbf{s}}) \sim p(\mathbf{y}|\tilde{\mathbf{s}})$, which produces detections $\mathbf{y} \in \mathcal{Y}$ from salient variables $\tilde{\mathbf{s}}$.

The planner in our case is deterministic, so

$$p(\mathbf{z}_t^1|\mathbf{y}) = \mathbb{1}_{\{\mathbf{z}_t^1=g(\mathbf{y})\}} \quad (6)$$

with the deterministic planning function $g(\mathbf{y})$, which can be used to rewrite Equation 5 as

$$\text{P-KL} = - \sum_{t=1}^{t=T} \log \left(\int p(\mathbf{z}_t^1|\mathbf{y})p(\mathbf{y}|\tilde{\mathbf{s}})d\mathbf{y} \right) \Big|_{\mathbf{z}^1=g(f(\mathbf{x}))}. \quad (7)$$

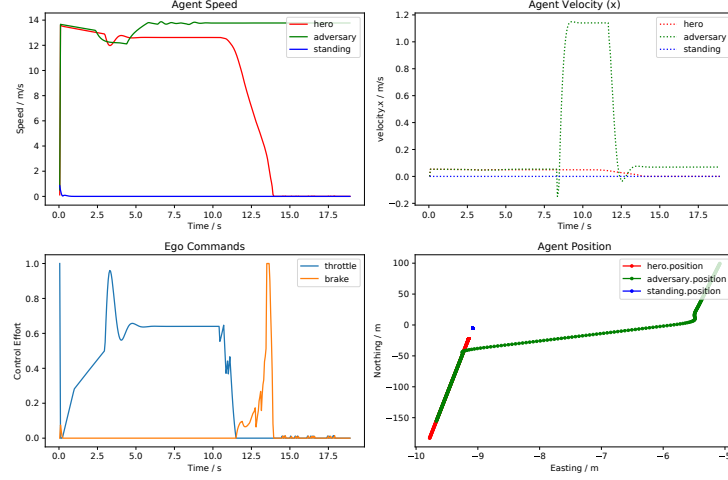


Figure 10: Diagnostics for simulation with the upstream detector outputs substituted for values generated by a neural network surrogate model.

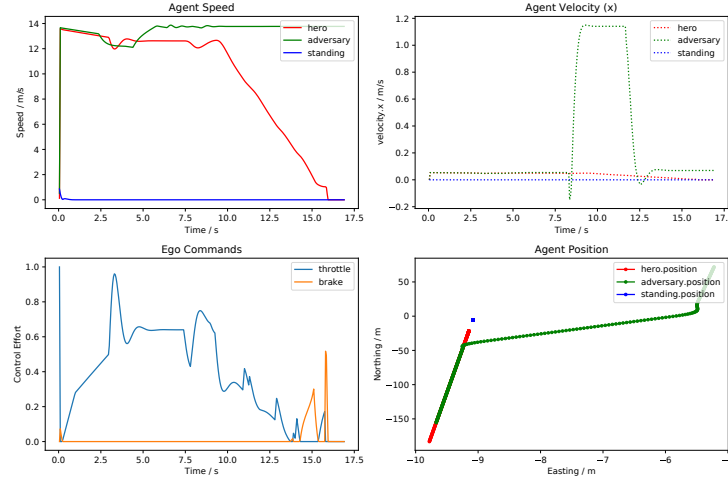


Figure 11: Diagnostics for simulation with the upstream detector outputs substituted for values generated by a logistic regression surrogate model.

Approximating the integral $\int p(\mathbf{z}_t^1 | \mathbf{y}) p(\mathbf{y} | \tilde{\mathbf{s}}) d\mathbf{y}$ with a kernel density estimator with bandwidth h obtained by sampling n times from $p(\mathbf{y} | \tilde{\mathbf{s}})$ yields

$$\text{P-KL} = - \sum_{t=1}^{t=T} \log \left(\frac{1}{nh} \sum_{\mathbf{y}} K \left(\frac{\mathbf{z}_t^1 - g(\mathbf{y})}{h} \right) \right) \Bigg|_{\mathbf{z}^1 = g(f(\mathbf{x}))}, \quad (8)$$

after which Jensen's inequality can be applied to obtain

$$\text{P-KL} \leq - \frac{1}{n} \sum_{t=1}^{t=T} \sum_{\mathbf{y}} \log \left(\frac{1}{h} K \left(\frac{\mathbf{z}_t^1 - g(\mathbf{y})}{h} \right) \right) \Bigg|_{\mathbf{z}^1 = g(f(\mathbf{x}))}, \quad (9)$$

which with a Gaussian kernel is equivalent to

$$\text{P-KL} \leq \frac{1}{n} \sum_{t=1}^{t=T} \sum_{\mathbf{y}} \left(\frac{(\mathbf{z}_t^1 - g(\mathbf{y}))^2}{2h^2} + \log(\sqrt{2\pi}h) \right) \Bigg|_{\mathbf{z}^1 = g(f(\mathbf{x}))}, \quad (10)$$

which is equal to the mean Euclidean norm metric (Eqn. 3) up to a constant and scaling factor, when the surrogate model is sampled repeatedly and the average taken.

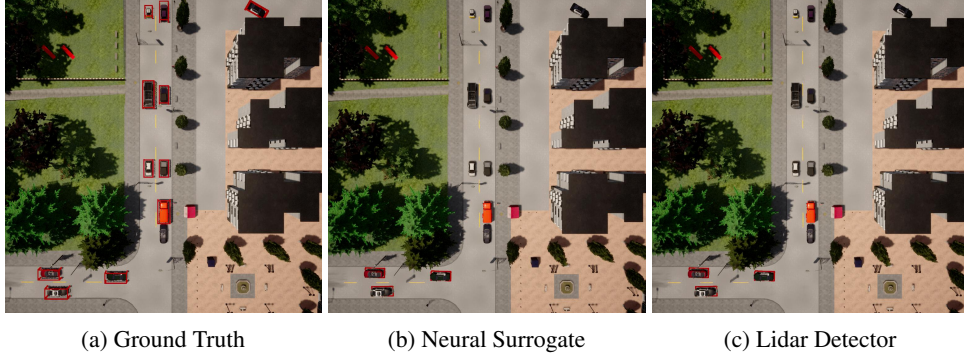


Figure 12: Example of vehicle false negative detection causing collision in Route 2. Detected objects are shown in red bounding boxes. Ego is the black vehicle colliding with the rear of the red van.

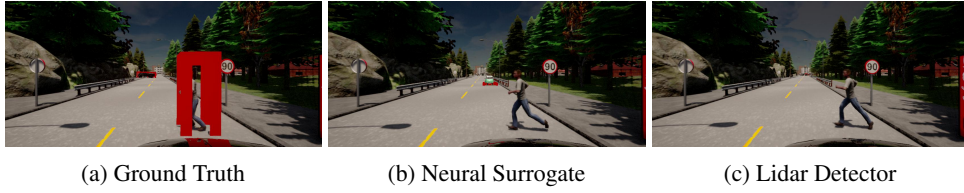


Figure 13: Example of pedestrian false negative detection causing collision in Route 1. Detected objects are shown in red bounding boxes.

H Example detection errors from simulation

Figures 12, 13 and 14 show examples of lidar detector errors which cause collisions which are reproduced by the surrogate model, and would not be reproduced if ground truth values were used in place of the backbone model in simulation.

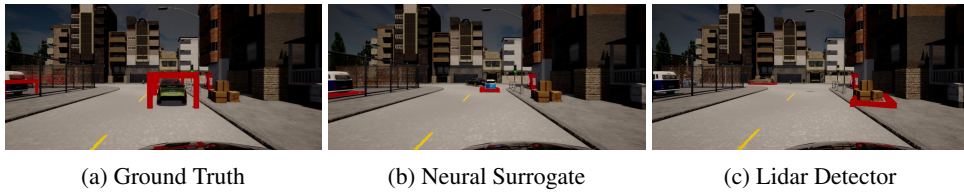


Figure 14: Example of false positive detection with a corner of the detected object in Ego's lane causing ego to stop in Route 0. Detected objects are shown in red bounding boxes.