

Quantifying Ghost Object Detectability in V2X Networks via Data-Driven Traffic Generation

Original

Quantifying Ghost Object Detectability in V2X Networks via Data-Driven Traffic Generation / Perrone, G., Bassi, F., Rapelli, M., Casetti, C.E.. - (In corso di stampa). (28th International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2026) Paris, France October 26–30, 2026).

Availability:

This version is available at: 11583/3016326 since: 2026-10-05T10:48:38Z

Publisher:

IEEE

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©9999 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Quantifying Ghost Object Detectability in V2X Networks via Data-Driven Traffic Generation

Giuseppe Perrone

Department of Control and
Computer Engineering
Politecnico di Torino
Torino, Italy
giuseppe.perrone@polito.it

Francesca Bassi

IRT SystemX
Palaiseau, France
francesca.bassi@irt-systemx.fr

Marco Rapelli

Department of Control and
Computer Engineering
Politecnico di Torino
Torino, Italy
marco.rapelli@polito.it

Claudio Ettore Casetti

Department of Control and
Computer Engineering
Politecnico di Torino
Torino, Italy
claudio.casetti@polito.it

Abstract—Vehicle-to-everything (V2X) communication enables connected vehicles to share sensor data, extending awareness beyond the individual field of view. This cooperation, however, exposes participants to the risk of receiving unreliable data, which is mitigated by misbehaviour detection mechanisms. A representative threat is the ghost object: a vehicle whose presence is announced in V2X messages but that does not physically exist on the road. Ghost objects may originate from deliberate attacks, e.g., a malicious participant claiming priority at an intersection, but also from perception faults on the sender’s side, such as partial occlusions causing duplicate tracks. Since misbehaviour detection relies on cross-checking observations from multiple participants, a ghost attack can be exposed only if honest, sensor-equipped vehicles actually cover its alleged location. Evaluating how often real traffic satisfies this condition is challenging: hand-crafted scenes lack statistical significance, while large-scale simulations obscure the conditions driving to detection failure.

We address this challenge by using a learned generative traffic model to synthesize large ensembles of realistic scenes (1,000 per experimental condition) over highway and intersection layouts. Through explicit sensor and occlusion modeling, we estimate the probability that a ghost object is observed by too few equipped vehicles to be reliably contradicted, thus remaining effectively undetected. We analyze how this probability varies with road topology, traffic density, environmental occlusion, and sensor penetration rate, under both awareness-only messaging (CAMs) and cooperative perception sharing (CPMs). We show that the dominant risk factor is the local visibility structure shaped by static obstructions preventing line-of-sight between equipped vehicles and the ghost location, and that the vulnerability transition shifts with penetration rate. Our analysis shows how generative traffic models may help to achieve a finer characterization of the performance of misbehaviour detection algorithms and ultimately a better understanding of the underlying phenomena.

Index Terms—Generative AI; Connected and Autonomous Vehicles; Cooperative Perception; Misbehaviour Detection

I. INTRODUCTION

Vehicle-to-Everything (V2X) communications allow vehicles to exchange information with low latency, and form the backbone of Cooperative Intelligent Transport Systems (C-ITS). In Europe awareness services rely mainly on two messages broadcast periodically: the Cooperative Awareness Message (CAM) [1], which each vehicle uses to advertise its own position and kinematic state; and the Cooperative Perception Message (CPM) [2], which disseminates the position

and motion data of nearby objects as detected by the sender’s onboard sensors. Together, they enable vehicles to build a picture of the environment that extends well beyond what any individual vehicle can perceive on its own.

This extended awareness, however, is only as trustworthy as the data it is built on. C-ITS address sources’ authentication and V2X message integrity through a vehicular Public Key Infrastructure (PKI). However, these mechanisms alone are insufficient to guard against the dissemination of unreliable kinematic data in CAM and CPM messages: a node may hold valid PKI credentials and yet broadcast false or distorted CAM and CPM data, whether due to sensor noise, hardware faults, or deliberate manipulation. *Misbehaviour detection* refers to the set of techniques used by the receiver to identify kinematic data that may be unreliable. These approaches [3] typically operate by exposing implausible values or by identifying inconsistencies across repeated observations over time or across multiple nodes. Characterizing the performance and effectiveness of this class of algorithms, however, is a challenging task.

We use the case of *ghost objects* as a running example throughout this work. We define a ghost object as a vehicle whose presence is declared in a CAM or CPM, but that has no actual physical existence on the road. Ghost objects may arise from malicious attacks, such as Sybil attacks carried out through CAMs [4]. They can appear from deliberate manipulation of CPMs by a malicious node seeking to distort the perceived environment for its own benefit, for instance by inducing surrounding vehicles to brake in order to claim right of way at an intersection [5]. They may as well stem from perception malfunctions on the sender’s side, such as partial occlusion causing a single physical object to be tracked as two distinct entities [6], both reported in CPMs. Regardless of its origin, a ghost object can be exposed by misbehaviour detection through cross-checking of the kinematic data reported by different participants. The detection success rests on an implicit assumption: that enough honest, CPM-capable vehicles can actually observe the location of the alleged object.

Assessing how often real traffic conditions satisfy this assumption is therefore a prerequisite for performance characterization. Yet, current validation practices through simulation cannot answer it: small sets of hand-crafted scenarios lack

statistical significance, while large-scale traffic simulations make it difficult to isolate the effect of individual factors, such as traffic density or the position of a ghost in a blind spot. In this work, we leverage a learned generative traffic model to try to bridge this gap. The generative model, trained on real-world driving data, provides the ability to populate a chosen road map with vehicles in realistic positions. This allows the sampling of thousands of distinct scenes under controlled conditions. We generate large ensembles of realistic scenes on chosen road topologies (highway and intersection); we evaluate, through explicit sensor and occlusion modeling, the probability that the presence of a ghost object cannot be reliably contradicted as a function of traffic density, environmental occlusion, and penetration rate of C-ITS capabilities. Thanks to this, we are able to statistically characterize in this study the *structural* exposure to ghost objects undetectable by misbehaviour detection algorithms.

The remainder of the paper is organized as follows. Section II reviews related work on cooperative perception security and data-driven scene generation. The scene generation framework and the visibility analysis pipeline are described in Sections III and IV, respectively. Experimental results are presented in Section V, and conclusions are drawn in Section VI.

II. RELATED WORK

A. Misbehaviour detection for Cooperative Perception

In Europe, the exchange of awareness and perception data is standardized by ETSI through the CAM [1] and CPM [2] specifications. Since PKI-based authentication cannot prevent a node with good credentials from broadcasting false kinematic data, misbehaviour detection techniques complement it by exposing implausible values and by checking the consistency of data reported by different nodes [3]. CPM-specific detectors extend these checks to perception data: they verify, e.g., that reported objects lie within the transmitter’s declared sensing area, and that objects jointly observable by multiple senders appear in all of their CPMs [7]. This cross-checking of independent observations is the mechanism that enables the detection of ghost objects [8]. The main limitation of this body of work lies in its evaluation. Attacks and defenses are typically validated on a handful of manually designed scenarios [5]. Large-scale simulations [4] and reference datasets such as VeReMi [9] only partially address the problem of characterizing the performance of these algorithms: they enable a fair comparison of detection algorithms, but no existing benchmark can quantify, e.g., how likely a fabricated object is to remain undetected, as a function of road geometry, traffic density, environmental occlusion, and number of sensor-equipped vehicles.

B. Data-Driven Traffic Scene Generation

In parallel, learned generative models have made it possible to synthesize realistic traffic scenes directly from real-world driving data, from autoregressive map-conditioned placement in SceneGen [10] and joint state-and-trajectory generation

in TrafficGen [11], to more controllable approaches such as Scenario Diffusion, UniGen, and DragTraffic [12], [13], [14]. These models provide exactly the missing ingredient identified above, i.e., the ability to sample arbitrarily many realistic scenes under controlled conditions. However, all these contributions focus on the generative tool itself, i.e., architecture, training objective, and controllability, and have not been applied to the security analysis of cooperative perception.

This work connects the two threads: we use a learned generative model to sample large, statistically meaningful ensembles of realistic scenes under controlled conditions, and we evaluate within each scene, through explicit sensor and occlusion modeling, the probability that a ghost target remains effectively unobserved under partial CPM penetration. To the best of our knowledge, no prior work combines learned scene generation, explicit visibility analysis, and sensor penetration rate modeling (i.e., the fraction of vehicles equipped with perception sensors) for this purpose.

III. GENERATIVE FRAMEWORK

The scene generator used in this study extends TRACE [15], a model from our previous work that synthesizes static traffic scenes conditioned on a description of the desired configuration. TRACE encodes the road geometry (lane boundaries, sidewalks, and other map features), expressed as a set of polylines, through a Transformer encoder, while a second encoder maps the scene description to a conditioning representation. A decoder attends jointly to both and samples, for each vehicle, its position, heading, and speed. The model is trained on the Waymo Open Motion Dataset [16] and produces scenes whose spatial statistics match those of real traffic. With respect to the published architecture, the variant employed here introduces two modifications aimed at improving controllability and sample diversity, both required to generate large, uniformly covering ensembles under precise experimental conditions. First, the original natural-language conditioning is replaced by a structured one: the space around a reference point is partitioned into six angular sectors, and for each sector the conditioning vector specifies the number of vehicles, the distance to the nearest one, and its speed. The original decoder constrains the output distribution to a fixed number of modes, and in practice tends to concentrate probability mass on a few dominant configurations, suppressing sample diversity. It is replaced by a Flow Matching decoder [17], which generates vehicle states by gradually transforming samples from an initial Gaussian noise distribution into plausible positions, headings, and speeds. Unlike the previous decoder, this formulation does not force the model to represent traffic scenes through a fixed set of output modes, and therefore allows more diverse vehicle configurations. Training uses minibatch Optimal Transport assignment [18], for straighter probability paths and more stable optimization. An auxiliary objective complements standard generation: in 25% of the training samples, a random subset of vehicles is injected into the decoder with its ground-truth states, and the model is trained to place the remaining vehicles consistently with the given context.




Fig. 1. Maps used in the visibility analysis. Panel (a) shows the highway scenario, panel (b) shows the four-way intersection without environmental blocking, and panel (c) shows the same intersection with blocking areas highlighted in blue. Lane centers, road edges, and crosswalk/road markings are shown in black, gray, and orange, respectively.

TABLE I
COMPARISON BETWEEN REAL AND GENERATED SCENE STATISTICS. W_1
DENOTES THE WASSERSTEIN-1 DISTANCE BETWEEN THE TWO
DISTRIBUTIONS.

Statistic	Unit	Real	Generated	W_1
Nearest-neighbour distance	m	8.52	8.63	0.22
Heading offset vs. lane	deg	6.36	5.68	0.76
Lateral distance to lane	m	0.72	0.72	0.17
Radial distance from origin	m	24.37	23.90	0.56
Speed	m/s	4.44	4.38	0.09
Collision rate	–	0.0014	0.0001	–
Off-road rate	–	0.317	0.188	–

This encourages each generated vehicle to be conditioned on its neighbors, improving the local consistency of spacing, heading alignment, and speed. To ensure that the generated scenes are suitable for visibility analysis, we compare 9,891 real scenes from the Waymo Open Motion Dataset (containing 117,476 vehicles) against 9891 generated scenes.

Table I summarizes the comparison over all evaluated statistics in terms of the Wasserstein-1 distance W_1 , a measure of the discrepancy between two probability distributions. All discrepancies are small with respect to the physical scale of the corresponding quantity (e.g., $W_1 = 0.22$ m on nearest-neighbour distances averaging 8.5 m); the largest relative one, on the lateral distance to lane center, is of the order of 15 cm and thus negligible for occlusion analysis between 2.0 m wide bounding boxes. Colliding bounding boxes are rarer in the generated scenes than in the real ones (0.01% vs. 0.14%), and their off-road rate (18.8%) is lower than that of the real data (31.7%), where parked vehicles frequently lie outside the driving corridor. Finally, the conditioning is respected: the number of generated vehicles per scene deviates from the requested one by 0.79 vehicles on average (mean absolute error). We conclude that the generator reproduces the spatial structure of real traffic with sufficient fidelity for the analysis that follows.

IV. METHODOLOGY

Misbehaviour detection for ghost objects relies on the assumption that independent participants can cross-check an announced vehicle’s presence. This section quantifies how often realistic traffic violates this assumption, with respect to the scene’s parametrization. We estimate this vulnerability through the following procedure: *i*) the experimental condition is fixed, namely the road topology and the number of vehicles n ; *ii*) for each experimental condition, $M = 1,000$ independent scenes are generated by resampling the model’s conditioning inputs (sector allocation, distances, and speeds) at each draw, so that each scene reflects a different plausible traffic configuration; *iii*) within each of the resulting scenes, the role of the ghost is assigned and the exposure risk is evaluated across a range of sensor penetration rates, each corresponding to a different random subset of vehicles being treated as non-equipped, i.e., blind to their surroundings. Verification works differently in the two V2X regimes, since they differ in whether a direct observation is shared with the rest of the network:

- in the CAM-only regime, an equipped vehicle that has the ghost’s position within its unobstructed sensor coverage resolves the doubt only for itself, since no perception data is exchanged; every other receiver remains unable to verify the claim on its own, regardless of what this one vehicle has seen;
- in the CAM+CPM regime, such observations are instead broadcast, so verification is pooled: once at least two independent equipped vehicles report the position as covered by their sensors, this corroborated evidence becomes available to every receiver, including those that cannot observe the location themselves.

For each experimental condition, we define the exposure risk as the fraction of generated cases in which the ghost object cannot be reliably contradicted by the network. Transmission and reception of V2X messages are assumed ideal, i.e., transmitted messages are received by all nodes within 100 m, and sensing is geometric and noise-free. Under these assumptions, the ability to contradict a ghost object depends only on whether enough sensors-equipped vehicles have line of sight to its

Selected Observer sensor-range plot (n=10, scene=2, penetration=0.5, map range=40 m)




Fig. 2. Example generated intersection scene with 10 vehicles and sensor penetration rate equal to 0.5. Sensor coverage is shown only for vehicle V2. The corresponding complete directed visibility matrix is reported on the right, where entries for non equipped vehicles are highlighted in grey.

alleged location. Communication losses and/or sensor noise would only make this condition harder to satisfy; therefore, the reported risk should be interpreted as a conservative lower bound.

A. Scene Generation and Road Configurations

We generate scenes on two road maps: a highway stretch and a four-way intersection, both taken from the Waymo Open Motion Dataset. To model environmental occlusion, the intersection map is also evaluated in an *obstructed* variant, in which static building footprints are added at the four corners of the intersection, alongside the road edges; these obstructions prevent vehicles traveling along one road from seeing what happens on the perpendicular arms until they are close to the intersection center, reproducing the occlusion conditions typical of built-up urban intersections. The map layouts are shown in Figure 1. For each map, we query the model for an increasing number of vehicles, from 4 to 25, requesting them uniformly across the map; this vehicle count n , together with the scenario type, defines each experimental condition. For the intersection, we additionally consider a main-direction variant, in which 80% of the requested vehicles are placed along a single arm, simulating one traffic flow dominating the intersection at a given moment, while the rest are spread over the other arms.

B. Sensor Configuration and Scene Representation

Once the scenes are generated, each vehicle is modeled as a rigid body with a fixed set of idealized perception sensors. Its physical footprint is an oriented bounding box of length 4.5m and width 2.0m, and its state is defined by position, heading, and speed in the scene reference frame. Each sensor is characterized by four parameters: mounting position in the vehicle reference frame (x, y) , orientation relative to the vehicle heading (Yaw), horizontal field of view (FOV), and maximum detection range (Range). In the considered scenario, each sensors-equipped vehicle is provided with one front long-range radar (FLRR), four short-range radars mounted at the

TABLE II
SENSOR PARAMETERS USED IN THE FULL SENSOR PRESET. MOUNTING OFFSETS ARE EXPRESSED IN THE VEHICLE LOCAL FRAME, WHERE POSITIVE x POINTS FORWARD AND POSITIVE y POINTS LEFT.

Sensor	x [m]	y [m]	Yaw [deg]	FOV [deg]	Range [m]
FLRR	2.25	0.00	0	24	250
Front-left SRR	2.25	1.00	45	150	30
Front-right SRR	2.25	-1.00	-45	150	30
Rear-left SRR	-2.25	1.00	135	150	30
Rear-right SRR	-2.25	-1.00	-135	150	30
FFC	0.50	0.00	0	50	150
360° LiDAR	0.00	0.00	0	360	100

vehicle corners (SRR), one front-facing camera (FFC), and one roof-mounted 360° LiDAR. The long-range radar and camera provide forward perception, the short-range radars cover the near-field regions around the vehicle, and the LiDAR provides omnidirectional coverage up to its maximum range. The parameters used for each sensor are shown in Table II.

Given a vehicle pose (x_i, y_i, θ_i) , a sensor with local mounting offset (x_s, y_s) and relative orientation α_s is transformed into the world frame as

$$\mathbf{p}_{i,s} = \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \mathbf{R}(\theta_i) \begin{bmatrix} x_s \\ y_s \end{bmatrix}, \quad \psi_{i,s} = \theta_i + \alpha_s,$$

where $\mathbf{p}_{i,s}$ is the position of sensor s mounted on vehicle i in the world frame, $\psi_{i,s}$ is the corresponding sensor viewing direction, and $\mathbf{R}(\theta_i) = \begin{bmatrix} \cos \theta_i & -\sin \theta_i \\ \sin \theta_i & \cos \theta_i \end{bmatrix}$ is the two-dimensional rotation matrix that maps vectors from the vehicle local frame to the world frame. We assume that a target is detected by a vehicle if it is visible to at least one of the mounted sensors.

C. Visibility Analysis

For each generated scene, visibility is evaluated independently for every ordered pair of vehicles and for every sensor mounted on the observing vehicle. Given an observer vehicle i , a sensor s , and a candidate target vehicle j , three conditions must hold simultaneously for the target to be detected.

1) *Range*: The distance from the sensor origin to the nearest point of the target's bounding box must not exceed the sensor's maximum range.

2) *Field of View*: The target's bounding box subtends an angular sector as seen from the sensor, bounded by its two outermost vertices. This sector is intersected with the sensor's angular field of view. If the intersection is empty — i.e., the entire bounding box lies outside the sensor's opening angle — the target is not detected.

3) *Non-blocking*: If the target passes the first two checks, we test whether any intermediate vehicle blocks the line of sight. A set of N rays (9 by default) is cast from the sensor origin, evenly distributed within the field-of-view-clipped angular sector. Only rays that intersect the target bounding box within the sensor's maximum range are considered valid. A valid ray is classified as blocked if it intersects the bounding box of any third vehicle at a strictly shorter distance (within a 10 cm tolerance). The target is declared occluded — and therefore not detected — only if every valid ray is blocked. In the constrained-intersection setting, the blue regions in Figure 1(c) are treated as roadside buildings or infrastructure. A ray is therefore marked as blocked if it intersects one of these regions before reaching the target. Detection is confirmed as soon as at least one valid ray reaches the target unobstructed.

By iterating this procedure over all sensors and all ordered observer–target pairs, we obtain the visibility matrix $\mathbf{V}$ associated with the generated scene. The directed entry $V_{i,j}$ is set to one if at least one sensor mounted on vehicle i detects vehicle j , and to zero otherwise. Diagonal entries are excluded by construction, as a vehicle cannot observe itself. An example of the resulting visibility matrix is shown in Figure 2.

D. Partial Sensor Deployment Modeling

Not all vehicles in a scene are assumed to carry perception sensors and C-ITS perception capabilities. To model this, a binary mask is applied to the visibility matrix after generation: given a scene of N vehicles, the penetration rate $p = (N - m)/N$ denotes the fraction of sensor-equipped vehicles, where m is the number of non-equipped ones. Non-equipped vehicles are treated as blind observers, i.e., their rows in the visibility matrix are masked, thus set to zero, as if they detect nobody.

Since the number of possible mask configurations grows combinatorially with scene size, for each generated scene a set of T masks is sampled at random, where T is chosen proportional to $\log_2(\binom{N}{m})$. Each mask is then applied independently to the visibility matrix of that scene, and the risk metrics are computed for each masked instance and averaged across the full set of sampled masks and generated scenes.

E. Risk Metrics

Two complementary metrics are computed for each masked scene. They quantify, at different levels of granularity, the structural exposure to ghost objects.

1) *CAM-only exposure risk*: The CAM metric serves as a baseline: it measures the exposure risk to ghost object when vehicles exchange only CAM, with no perception data transmitted.

Since no perception data is shared, each receiver must assess every announced vehicle on its own: any vehicle in the scene is

therefore treated as a potential ghost, and the metric measures the average fraction of observer–target pairs in which the observer cannot verify the target's existence. Given a masked visibility matrix $\tilde{\mathbf{V}}$, the unseen fraction for a single scene and mask instance is

$$u = 1 - \frac{\sum_{i \neq j} \tilde{V}_{i,j}}{N(N-1)}.$$

The reported CAM value is averaged over all evaluated masks and over all retained generated scenes for the corresponding vehicle count.

2) *CPM exposure risk*: The CPM metric captures the exposure once vehicles also share Cooperative Perception Messages, i.e., they broadcast what their sensors perceive to all other participants.

In this regime, observations are pooled across all equipped vehicles, so the relevant question is no longer the average pair-wise visibility, but whether a specific position exists where a ghost could escape collective scrutiny. For each masked scene, the ghost candidate is therefore identified as the vehicle seen by the fewest equipped observers, i.e., the worst-case position whose fabricated presence would be hardest to contradict from the shared perception picture.

Denoting by g the number of equipped vehicles that detect this target, a ghost object is considered undetected if $g < k$, with detection threshold $k = 2$: a single witness cannot be distinguished from a faulty or malicious node, so at least two concordant, independent observations are required before the network can reliably assert or refute the vehicle's presence. The sensitivity of the results to this choice is assessed in Section V. Since direct observers can directly recognize the false-vehicle as a ghost, the fooled fraction for a single masked scene is:

$$f = \begin{cases} \frac{N-1-g}{N-1}, & \text{if } g < 2, \\ 0, & \text{otherwise.} \end{cases}$$

The reported CPM value is averaged over all sampled masks and generated scenes.

V. RESULTS

This section presents a quantitative analysis of the exposure to ghost objects across the four experimental conditions defined in Section IV, i.e., the highway stretch, the four-way intersection with uniform and main-direction traffic, and its obstructed variant, illustrated in Figure 1. Together, these conditions allow us to quantify the exposure as a function of traffic density, spatial arrangement, and environmental occlusion.

A. CAM-only Exposure to Ghost attack

Figure 3 shows the CAM-only exposure risk as a function of the penetration rate, for different numbers of vehicles. Increasing the penetration rate always reduces the exposure, but the improvement is gradual and approximately linear. The qualitative behavior is remarkably similar across all scenarios. The obstructed intersection does not exhibit a fundamentally different trend; instead, it consistently presents higher exposure risk values because the obstacles remove visibility links




Fig. 3. CAM-only exposure risk as a function of the penetration rate. Each curve corresponds to a different number of vehicles, indicated by the colorbar.




Fig. 4. CAM-only exposure risk as a function of the number of vehicles in the scene. Each panel corresponds to one scenario, and each curve is colored by the penetration rate, i.e. the percentage of vehicles equipped with sensors.

from the sensing matrix, creating additional blind areas that cannot be resolved through CAM messages alone.

The effect of traffic density is shown in Figure 4. Exposure generally grows with the number of vehicles, as mutual occlusion increases with density. Unlike the penetration rate, however, this effect is nonlinear and scenario-dependent, with curve shapes that vary with the road layout and visibility conditions. The two variables also interact: the influence of density is pronounced at high penetration rates, whereas below a penetration rate of approximately 30% the curves flatten and the exposure becomes nearly independent of the number of vehicles in the scene.

B. CPM Exposure Risk

Figures 5 and 6 report the CPM exposure risk as a function of the number of vehicles and of the penetration rate, respectively. The role of traffic density is reversed with respect to the CAM case: each additional equipped vehicle is a potential observer of the ghost location, and this effect dominates the increased mutual occlusion, so that exposure decreases as the scene becomes denser. The penetration rate drives a sharp transition instead: at low rates the exposure remains high




Fig. 5. CPM-based exposure risk as a function of the penetration rate. Each panel corresponds to one scenario, and each curve is colored according to the total number of vehicles present in the scene.




Fig. 6. CPM-based exposure risk as a function of the number of vehicles in the scene. Each panel corresponds to one scenario, and each curve is colored according to the penetration rate, i.e. the percentage of vehicles equipped with sensors and capable of sharing perception information.

regardless of density, since few vehicles contribute to the collective perception process, while as the rate grows the curves rapidly drop toward zero. The transition becomes steeper as the number of vehicles increases, because in dense scenes a small increment in penetration rate is sufficient to bring many new observers online. This behavior is consistent across all scenarios. The obstructed intersection is the worst case and exhibits the shallowest transition; crucially, the obstacles do not alter the underlying mechanism, but act as an offset and a scaling factor, shifting the transition toward higher penetration rates and reducing its steepness. Road topology and traffic distribution have, by contrast, a comparatively minor influence. To characterize the shape of the vulnerability transition across scenarios and vehicle counts, we extract three scalar descriptors from each CPM exposure curve of Figure 5. The onset threshold is the highest penetration rate at which the CPM risk first exceeds 1%, i.e., the point at which a non-negligible fraction of scenes become exploitable as the fraction of equipped vehicles decreases from 100%. The transition width is the span of penetration rate between the 10th and 90th




Fig. 7. Exposure Risk in case of CPM transmission, transition characteristics as a function of the number of vehicles in the scene, evaluated at maximum penetration rate. Each panel corresponds to one scenario. Left: onset threshold (masked-vehicle fraction at which CPM risk first exceeds 1%). Center: transition width. Right: maximum slope. Curves are colored by scenario type.

percentiles of the total risk rise for a given curve, measuring how abruptly the network switches from near-zero to high exposure. The maximum slope is the peak local derivative of the curve, expressed in percentage points of exposure risk per percentage point of penetration rate decrease, and captures the steepness of that switch.

Figure 7 reports these three descriptors as a function of the number of vehicles n , evaluated at maximum penetration (100% equipped). All three metrics confirm the picture drawn from Figures 5 and 6. The open intersection, highway, and main-direction scenarios cluster tightly in all three panels: their onset thresholds are nearly identical, their transition widths shrink at the same rate as n grows, and their maximum slopes increase in parallel, confirming that road topology and traffic distribution have only a minor effect on the shape of the transition. The obstructed intersection is displaced consistently toward higher onset values and wider, shallower transitions: obstacles do not create a qualitatively different regime, but act as an offset that demands a higher penetration rate before the network reaches safety, and a scaling factor that makes the residual transition more gradual.

Figure 8 repeats the CPM analysis while varying the detection threshold k , i.e., the minimum number of concordant equipped observers required to contradict the ghost ($k = 2$ elsewhere in this work). Increasing k shifts the vulnerability transition toward higher penetration rates and makes it more gradual, since more independent witnesses must simultaneously cover the ghost location; relaxing it to $k = 1$, i.e., trusting any single observer, produces the opposite effect. Crucially, the threshold acts in the same way as the environmental obstructions, i.e., as an offset and a scaling factor on the transition, without altering the qualitative behavior: for every k , exposure still decreases with density, the transition steepens in denser scenes, and the obstructed intersection remains

consistently the most exposed. The findings of this study are therefore robust to the specific choice of k , which acts as a tunable security margin rather than a structural assumption.

VI. CONCLUSION

In this work, we leveraged a learned generative traffic model to characterize the structural exposure of V2X networks to ghost objects. The key methodological contribution lies in the evaluation paradigm itself: by sampling 1,000 independent, realistic scenes per experimental condition under explicit control of road topology, traffic density, environmental occlusion, and sensor penetration rate, we disentangle the contribution of each individual factor and extract statistically robust relationships that neither hand-crafted scenarios nor large-scale traffic simulations can provide. Since our analysis abstracts the communication channel, sensor noise, and concrete detector logic, relaxing each of these can only increase the measured exposure, making our estimates conservative lower bounds.

The analysis yields three main findings. First, in the CAM-only regime, exposure decreases approximately linearly with the penetration rate and grows with traffic density due to mutual occlusion. Second, when CPM transmission is active, the role of density is reversed and the penetration rate drives a sharp transition from high vulnerability to near-zero risk; crucially, our framework allows one to directly read off the minimum penetration rate required to bring exposure below any target threshold, as a function of traffic density and road configuration. Third, the dominant risk factor is the local visibility structure shaped by static obstructions, which act as an offset and a scaling factor on the vulnerability transition, shifting it toward higher penetration rates and making it more gradual, while road topology and traffic distribution have comparatively minor influence.

On the modeling side, cooperative perception risk models should explicitly account for static occlusions and line-of-




Fig. 8. Sensitivity of the CPM exposure risk to the detection threshold k (a ghost is undetected if observed by fewer than k equipped vehicles), for the open (top) and obstructed (bottom) intersections. Curves are colored by the number of vehicles in the scene.

sight structure. On the evaluation side, our results demonstrate that data-driven scene generation can serve as a principled benchmarking substrate for V2X security. Future work will extend the framework to full trajectory generation — enabling temporal consistency checks and dynamic attack scenarios — and will couple it with realistic channel models and concrete detectors to quantify the gap between structural and operational exposure.

REFERENCES

- [1] ETSI, “Intelligent transport systems (ITS); facilities layer; cooperative awareness service; release 2,” European Telecommunications Standards Institute, Technical Specification ETSI TS 103 900 V2.2.1, Feb. 2025.
- [2] —, “Intelligent transport system (ITS); vehicular communications; basic set of applications; collective perception service; release 2,” European Telecommunications Standards Institute, Technical Specification ETSI TS 103 324 V2.1.1, Jun. 2023.
- [3] R. W. Van Der Heijden, S. Dietzel, T. Leinmüller, and F. Kargl, “Survey on misbehavior detection in cooperative intelligent transportation systems,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 779–811, 2018.
- [4] J. Kamel, M. R. Ansari, J. Petit, A. Kaiser, I. B. Jemaa, and P. Urien, “Simulation framework for misbehavior detection in vehicular networks,” *IEEE transactions on vehicular technology*, vol. 69, no. 6, pp. 6631–6643, 2020.
- [5] J. Zhang, I. B. Jemaa, and F. Nashashibi, “Trust management framework for misbehavior detection in collective perception services,” in *2022 17th International Conference on Control, Automation, Robotics and Vision (ICARCV)*. IEEE, 2022, pp. 596–603.
- [6] R. Dominguez, E. Onieva, J. Alonso, J. Villagra, and C. Gonzalez, “Lidar based perception solution for autonomous vehicles,” in *2011 11th International Conference on Intelligent Systems Design and Applications*. IEEE, 2011, pp. 790–795.
- [7] M. R. Ansari, J.-P. Monteuis, J. Petit, and C. Chen, “V2x misbehavior and collective perception service: Considerations for standardization,” in *2021 IEEE Conference on Standards for Communications and Networking (CSCN)*. IEEE, 2021, pp. 1–6.
- [8] X. Liu, L. Yang, I. Alvarez, K. Sivanesan, A. Merwaday, F. Oboril, C. Buerkle, M. Sastry, and L. G. Baltar, “Miso-v: Misbehavior detection for collective perception services in vehicular communications,” in *2021 IEEE intelligent vehicles symposium (IV)*. IEEE, 2021, pp. 369–376.
- [9] R. W. Van Der Heijden, T. Lukaseder, and F. Kargl, “Veremi: A dataset for comparable evaluation of misbehavior detection in vanets,” in *International conference on security and privacy in communication systems*. Springer, 2018, pp. 318–337.
- [10] S. Tan, K. Wong, S. Wang, S. Manivasagam, M. Ren, and R. Urtasun, “Scenegen: Learning to generate realistic traffic scenes,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 892–901.
- [11] L. Feng, Q. Li, Z. Peng, S. Tan, and B. Zhou, “Trafficgen: Learning to generate diverse and realistic traffic scenarios,” *arXiv preprint arXiv:2210.06609*, 2023.
- [12] E. Pronovost, M. R. Ganesina, N. Hendy, Z. Wang, A. Morales, K. Wang, and N. Roy, “Scenario diffusion: Controllable driving scenario generation with diffusion,” in *Advances in Neural Information Processing Systems*, 2023.
- [13] R. Mahjourian, R. Mu, V. Likhoshesterov, P. Mouglin, X. Huang, J. Mesias, and S. Whiteson, “Unigen: Unified modeling of initial agent states and trajectories for generating autonomous driving scenarios,” in *2024 IEEE International Conference on Robotics and Automation*, 2024.
- [14] S. Wang, G. Sun, F. Ma, T. Hu, Q. Qin, Y. Song, L. Zhu, and J. Liang, “Dragtraffic: Interactive and controllable traffic scene generation for autonomous driving,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024, pp. 14 241–14 247.
- [15] G. Perrone, M. Rapelli, and C. Casetti, “Trace — traffic autoregressive scene generation,” in *Proceedings of IEEE VTC2026-Fall*, 2026, accepted for publication.
- [16] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou *et al.*, “Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9710–9719.
- [17] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, “Flow matching for generative modeling,” *arXiv preprint arXiv:2210.02747*, 2023.
- [18] A. Tong, K. Fatras, N. Malkin, G. Huguet, Y. Zhang, J. Rector-Brooks, G. Wolf, and Y. Bengio, “Improving and generalizing flow-based generative models with minibatch optimal transport,” *Transactions on Machine Learning Research*, 2024.