A Statistical Characterization of the Actual Cooperative Perception Messages and a Generative Model to Reproduce Them

Mattia Andreani[1,2], *Student Member, IEEE*, Luca Lusvarghi[2,3], *Associate Member, IEEE* and Maria Luisa Merani[1,2], *Senior Member, IEEE*

[1]Department of Engineering "Enzo Ferrari", University of Modena and Reggio Emilia, Modena, 41125, Italy
[2] Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT), Parma, 43124, Italy
[3] UWICORE Laboratory, Universidad Miguel Hernández de Elche (UMH), Elche, 03202, Spain
(e-mail: 241125@studenti.unimore.it, llusvarghi@umh.es, marialuisa.merani@unimore.it)

***Abstract*—This paper provides two novel contributions to vehicular cooperative perception. Firstly, it puts forth an approach to generate the actual perception messages broadcasted by connected autonomous vehicles. Relying on data gathered by autonomous vehicles and originally collected for computer vision purposes, it produces perception messages in accordance with the standard ETSI rules. The statistical properties of the messages are determined, showing that their size is remarkably affected by the driving scenario and the policy adopted to discern when an object is seen by the vehicle, and to a lesser extent by the selection of the message generation frequency. Secondly, the paper proposes a generative model to synthetically replicate the sequences of perception messages. The ability of the model to successfully capture the characteristics and the temporal correlation of the real data is demonstrated in a reference scenario. The model adoption is promising in large-scale numerical simulations, where the perception messages of many vehicles have to be faithfully reproduced.**

***Index Terms*—Cooperative Perception Messages, Connected Autonomous Vehicles, CPM, CAV, TimeGAN**

I. Introduction

Communications play a fundamental role in the Intelligent Transport System (ITS) domain, as vehicles, Road Side Units (RSUs), and potentially all street "inhabitants" become involved in the process of information sharing. Initially, vehicles are expected to exchange messages carrying local status information, e.g., the vehicle's current speed and location, to support basic safety applications. In the second development stage, cooperative perception is foreseen, where vehicles also broadcast what they "see", to augment the individual perception and make autonomous driving a reality. This collective approach aims to satisfy the requests of challenging use cases, such as the active protection of non-connected Vulnerable Road Users (VRUs), e.g., pedestrians and bikers, and the planning of cooperative trajectories and maneuvers of Connected Autonomous Vehicles (CAVs), to name a few, meaningful examples.

Cooperative sensing is a field that necessitates thorough scrutiny to understand the relation between the wealth of information gathered by CAVs and the advantages, along with the cost, of sharing such information through the wireless channel. To build a situational awareness that is as complete as possible, Cooperative Perception Messages (CPMs) have been introduced by the European Telecommunications Standards Institute (ETSI) in the technical standard TS 103 324 V2.1.1 [1]. They convey a summary, in the form of a list of attributes, of the objects (road users and obstacles) the CAV detected with sufficient confidence. These messages are transmitted through point-to-multipoint communications by vehicles and also connected RSUs.

Among the papers investigating cooperative perception, [2] identified the ETSI rules in [1] as responsible for generating frequent messages, each providing information about a modest number of objects. Although significant, this work relied on simulation and assumed CAVs equipped with a single camera sensor. The outcomes of [3] quantified the usefulness of broadcasting cooperative perception messages at different penetration rates of CAVs, introducing the notion of perceived vehicular safety. Here too, the few vehicle sensors and their characteristics were simulated. A simulative approach was undertaken in [4] too, where CAVs equipped with two RADAR sensors were considered. This study compared alternative redundancy mitigation techniques, which are introduced to avoid excessive traffic loads on the radio channel. In [5], CPMs and the messages transmitted by VRUs were jointly examined, the CAVs being equipped as in [4]; by simulation, the paper determined the improvement in the rate of detection of VRUs and in the time required to identify their presence.

In reality, autonomous vehicles are provisioned with numerous cameras, LiDARs, RADARs, and high-definition maps, which augment the vehicle's perception via sensor fusion. A realistic characterization of the Autonomous Vehicle (AV) as a data source is fundamental, as it allows a sound estimate of the requirements set on the communication channel when the AV possesses the communication capability, that is, when it is a CAV. This is where the first contribution of the present work lies. Specifically, this study generates and statistically characterizes the CPMs conveying information about the ob-

jects seen by an actual AV. To achieve this goal, two annotated datasets collecting the scenes recorded by a fully-equipped AV driving in different environments [6], [7] are post-processed, and the CPMs the AV would transmit if it were a CAV are generated in accordance with the rules proposed by ETSI. Two simple detection policies that ascertain the objects perceived by the CAV are also put forth. The policies identify two limit cases and realistically quantify the range where the number of detected objects falls; as it is this number that primarily affects the CPM size, a true estimate of the interval of the latter is provided. Furthermore, the study reveals that:

- the message size primarily depends on the driving scenario and the object detection policy, and to a lesser extent on the message generation frequency;
- the number of objects the CAV detects significantly differs from the one previously published in the literature and estimated via simulative approaches.

To overcome the limitation in size of the set of available CPMs, the present work also proposes a specific neural network, the Time-series Generative Adversarial Network (TimeGAN), to obtain synthetic replicas of the CPMs exhibiting the same statistical properties as the original ones. The meaningful example of the highway is examined, demonstrating that the model is successful and in the future can be utilized in large-scale numerical simulations involving a remarkable number of vehicles, each generating CPMs.

The rest of the paper is organized as follows: Section II recalls the notion of cooperative perception service and details the format of CPMs; Section III illustrates the approach followed when generating CPMs from the actual datasets; Section IV introduces the TimeGAN-based model and the metrics employed to tune its parameters and evaluate its performance; Section V discusses the results and Section VI draws the conclusions.

II. Cooperative Perception and CPMs

Autonomous vehicles exploit different types of sensors to gather information about the surrounding environment. Cooperative perception services have been proposed to share this information with other vehicles and the roadside infrastructure, relying on Vehicle-to-Everything (V2X) communications. This allows the expansion of the Field of View (FoV) of the vehicle, which not only perceives the surrounding environment through its own equipment but becomes aware of what other vehicles and the RSUs see. The shared information is expected to be organized in CPMs, proposed in a dedicated ETSI standard [1].

A CPM comprises several fields, as illustrated in Fig. 1: an ITS-PDU header containing the protocol version, the message type, and the ID of the ITS station (vehicle or RSU) that generated the message; a management container reporting the transmitter type and reference position, along with optional segmentation information; a station data container including a container for the originating ITS station; a sensor information container reporting the ID, type, and detection area of each onboard sensor which the disseminating station is equipped

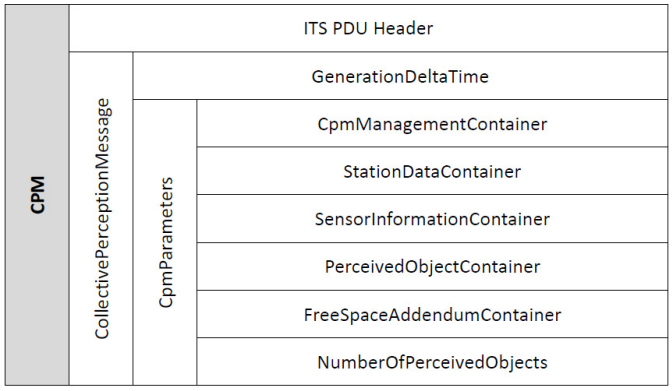

Fig. 1. CPM structure [1]

with; a perceived object container detailing the perceived objects by means of their ID, type, reference position and time of measurement, and a free space addendum container used to describe the free space areas within the sensor detection areas.

The ITS PDU header and the management container are mandatory, all the remaining containers are optional. Moreover, the sensor information container should be included in a CPM if the time elapsed since the last inclusion is greater than 1 s. Considering perceived objects, an object has to be included in a CPM if specific conditions are satisfied. In greater detail, if the object class does not correspond to either the person or animal class, the object inclusion rules are the following:

1) The object is new with respect to the last generated CPM;
2) The Euclidean absolute distance between the current position of the object and its position lastly included in a CPM exceeds 4 m;
3) The difference between the current absolute speed of the object and its speed lastly included in a CPM exceeds 0.5 m/s;
4) The difference between the current heading of the object and its heading lastly included in a CPM is 4°;
5) The time elapsed since the last time the object was included in a CPM is greater than 1 s.

On the other hand, if the object class corresponds to the person or animal class, the rules modify in:

1) The object is new with respect to the last generated CPM;
2) If the object list contains at least one object of class person or animal that has not been included in a CPM in the past 500 ms, all objects of class person or animal should be included in the currently generated CPM.

As regards the dissemination of such messages, the ETSI standard suggests that the minimum time elapsed between consecutive CPMs should be equal to or larger than T_{genCPM}, 100 ms $\leq T_{genCPM} \leq$ 1000 ms.

III. CPMs generated by real-world CAVs

A. The Examined Datasets

The first examined dataset is NuScenes [6] [8], whose scenes were sampled at a 2 Hz rate by an AV equipped

with six cameras, five RADARs, one LiDAR, one GPS, and an Inertial Measurement Unit (IMU), while the AV traveled in four different urban and suburban scenarios. NuScenes provides further information drawn from the onboard sensors, namely, the number of RADAR points from the ensemble of RADARs and the number of LiDAR points associated with each object, as well as the object visibility level from the ensemble of the six cameras. Four different visibility levels are present: $[0, 40\%]$, $[40\%, 60\%]$, $[60\%, 80\%]$, and $[80\%, 100\%]$.

The second dataset is Cirrus [7] [9], whose scenes were recorded by an AV in an urban and a highway scenario. The vehicle is equipped with one camera, one Gaussian LiDAR and one uniform LiDAR, two GPSs, and an IMU. Cirrus scenes are sampled at a 1 Hz frequency. Cirrus structure is simpler than NuScenes, as no information about the onboard sensors is reported. Namely, there is no indication about the number of LiDAR points associated with an object, nor an indication about the object visibility level. As a consequence, every object present in this dataset is assumed to be detected by the CAV with the highest confidence level.

Together, NuScenes and Cirrus recordings cover all the macroscopically relevant driving environments, i.e., urban, suburban, and highway.

B. CPM Generation and Fusion Rules

The previously described datasets were conceived for computer vision purposes, i.e., to train and test the effectiveness of neural network architectures in classifying the objects "seen" by the AV. We exploit them for a different purpose, that is, we process their elements to generate actual CPM sequences, that we also term traces. To this end, we observe that each scene is a sequence of annotated frames, where each annotation corresponds to an object and is associated with a bounding box. From the attributes of the bounding boxes, it is possible to understand whether they correspond to a person or an animal; furthermore, the dynamics of the bounding boxes allow us to determine whether the ETSI object inclusion rules of Section II are satisfied. In greater detail, examining how the bounding box of each object moves and modifies over time allows: (i) to spot objects that enter and disappear from the vehicle's viewing horizon; (ii) to identify the position of every object in the scenes. Concerning the last point, it is worth observing that the GPS and the IMU jointly provide the position of the barycenter of the bounding boxes in the global coordinate system. Hence, it is possible to compute the variation in distance and orientation undergone by every object that does not correspond to either the person or animal class, and then check whether the inclusion rules 2) and 4) of Section II are verified. The speed variation is necessary to verify the inclusion rule 3) and is computed under the hypothesis that the objects obey the uniform rectilinear motion, an acceptable assumption, as adjacent scenes differ by either 500 or 1000 ms.

For NuScenes, which provides a richer collection of elements about each object, we also propose two alternative object detection policies, termed relaxed and restricted. When the relaxed policy is adopted, an object is considered detected by the CAV if the object visibility level is greater than 40% *or* there is at least one RADAR point *or* there is at least one LiDAR point associated to the object bounding box. When the restricted policy is employed, an object is considered detected by the CAV if the object visibility level is greater than 40% *and* there is at least one RADAR point *and* there is at least one LiDAR point associated to the object bounding box. The adoption of either one of the two policies differently affects the number of objects perceived by the CAVs, on which the inclusion rules proposed by ETSI are next applied; hence, non-negligible changes in the CPM size are expected.

After applying the ETSI object inclusion rules to the perceived objects, the actual CPM traces are obtained by exploiting Vanetza [10], an open-source tool whose implementation adheres to the ASN1 syntax of [1]. Vanetza automatically computes the size in bytes of the CPMs, provided the number of objects to include in the message has been previously determined. The analysis of the produced traces allows to determine the traffic load the CAV places on the radio channel, shedding some light on the bandwidth requirements of CPM transmission.

IV. A Neural Network Model to Generate CPMs

For each driving environment, it would be highly desirable to have a model of the CPM traffic generated by the CAV. This would allow to perform large-scale numerical simulations where the size and temporal pattern of the perception messages are realistically re-created. A possible approach consists of resorting to an Artificial Intelligence (AI) generative model, properly trained through actual data. We have chosen to follow this path, concentrating on the number of objects N_{obj} the CAV perceives and includes in the CPMs, as the CPM size and the correlation between the size of consecutive CPMs are mainly determined by this number. The number N_{obj} evolves over time, i.e., it is a time series, which suggests the selection of a specific instance of a neural network, the so-called TimeGAN [11]. We resort to and properly train it, considering different combinations of its parameters, specifically, the length of the generated temporal sequences, the batch size adopted for the training procedure, and the number of training iterations.

In greater detail, the TimeGAN is composed of: (i) an embedding function; (ii) a recovery function; (iii) a sequence generator; (iv) a sequence discriminator. The embedding function serves the purpose of moving from the feature space to the latent space. The latter is a compressed representation of the input data that preserves its statistical properties. The outputs of the embedding function are named real latent codes. The recovery function allows to move in the opposite direction, i.e., from the latent space to the original feature space, to obtain the so-called reconstructions of the real latent codes. The sequence generator takes as input random vectors chosen from vector spaces over which known distributions are defined and produces outputs belonging to the latent space that are denoted as synthetic latent codes. Finally, the sequence discriminator

starts from latent codes, either real or synthetic, to produce classifications as outputs.

To assess the validity of this solution, as in [11] we compute: (i) the discriminative score, which quantitatively measures the similarity between the original data and the synthetic data generated by the TimeGAN; (ii) the predictive score, which measures the ability of the TimeGAN to capture the temporal correlation of the data.

We parallel these AI quantitative metrics with more traditional measures employed in probability theory; namely, we evaluate the Kullback-Leibler (KL) divergence, D_{KL}, to assess how close the Probability Mass Function (PMF) of the number of objects the CAV includes in the CPMs and the PMF of the number of objects in the artificial sequences are. This indicator is defined as

$$D_{KL}(p\|q) = \sum_{i\in I} P(i)\log_{10}\frac{P(i)}{Q(i)} \tag{1}$$

where $P(i)$ is the PMF computed from the actual data, $Q(i)$ the PMF computed from the synthetic sequences, and I is the sample space. Both PMFs are *a posteriori* determined.

We additionally evaluate and confront the Auto Correlation Function (ACF) of the real and artificial time series. The ACF, which is a temporal sequence, is computed as

$$r_k = \frac{\sum_{i=1}^{M}\sum_{n=0}^{N} x_n^{(i)} x_{n+k}^{(i)}}{s_0^2 \cdot M \times N}, \quad k = 0, 1, \ldots, N \tag{2}$$

where k denotes the temporal lag, M the number of time series, N the number of elements in each time series, $x^{(i)}(n)$, $n = 0, 1, \ldots, N$, the single time series, and s_0^2 is

$$s_0^2 = \frac{\sum_{i=1}^{M}\sum_{n=0}^{N}(x_n^{(i)} - \bar{x}^{(i)})^2}{M \times N}. \tag{3}$$

V. Numerical Results

A. Statistical Characterization of the CPMs

We first determine the PMF of the CPM size for both datasets. For NuScenes, we select two of the settings where the recordings took place, namely, Singapore Holland Village and Boston Seaport, the former being a quiet, residential neighborhood, the latter a busy touristic district. We take them as good references for a suburban and an urban environment, respectively. From now on, we refer to Singapore Holland Village as the suburban setting and to Boston Seaport as the urban one. Moreover, we denote by X the CPM size. The number of 20 second-long scenes recorded in the suburban setting is equal to 85, which amounts to 26 minutes; the scenes are 467 in the urban environment, for 156 minutes. We generate CPMs with a periodicity equal to $T_{genCPM} = 500$ ms and 1 s. These values are compliant with the ETSI standard and also coherent with the sampling rate of the dataset, which is 2 Hz.

For the suburban scenario, Fig. 2(a) reports the PMF of X when the relaxed policy is considered. Instead, Fig. 2(b) refers to the restricted policy. The comparison between the two figures shows the impact of the detection rule on the

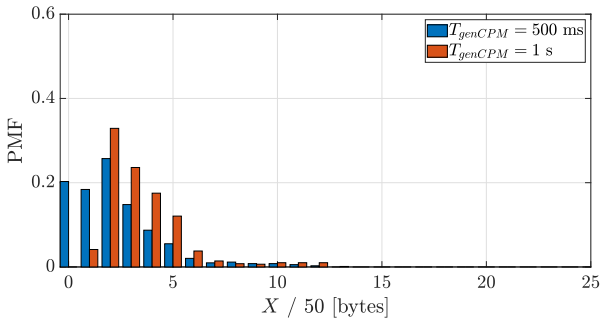

(a) Relaxed policy.

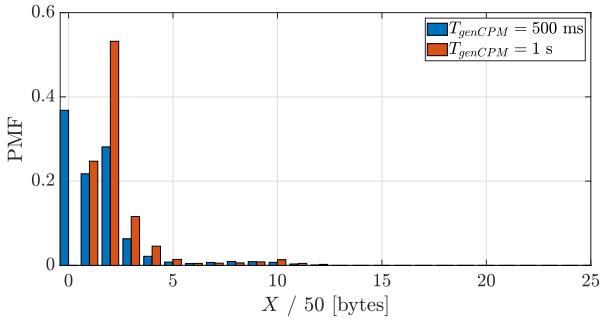

(b) Restricted policy.

Fig. 2. PMF of X, NuScenes dataset, suburban setting.

CPM size. Shorter messages are obtained when the restricted policy is employed, the reason being that an object is not included in subsequent messages when it is not perceived with a sufficiently high confidence level. The figures also show the T_{genCPM} impact on the PMF: if the CPM generation period is lower, smaller CPM sizes become more likely. As a matter of fact, a shorter interval between consecutive CPMs decreases the probability of observing a variation in the objects' position, speed, and heading, which in turn leads to fewer objects being included in the message.

For the urban scenario, the PMF of X is plotted in Figs. 3(a) and (b), for the two distinct fusion rules. The comparison with the former figures indicates that the setting has a notable influence on the PMF. Owing to the large number of objects populating the scenes in this touristic area (pedestrians, bicycles, and animals), the X range is wider and the PMF is shifted to the right with respect to the suburban scenario. These figures also confirm that the detection rule plays a non-negligible role in determining the actual location and shape of the PMF.

The statistical characterization of the CPM size in the urban and suburban scenarios provides a valuable insight into the communication requirements placed by each CAV on the radio channel. In this regard, Table I reports the average and the maximum CAV data rate for the NuScenes dataset considering the same detection policies, scenarios, and T_{genCPM} values as in Figs. 2 and 3. The average values in Table I reflect the CPM size trends. Namely, the relaxed policy exhibits the largest average data rates, as visible in the first and third columns. On the other hand, the maximum data rates capture

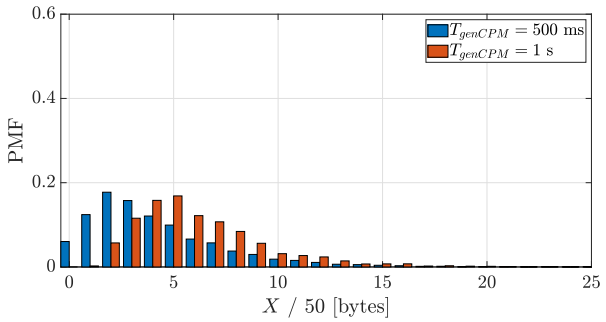

(a) Relaxed policy.

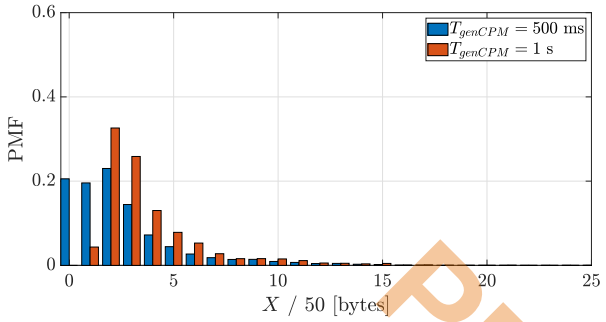

(b) Restricted policy.

Fig. 3. PMF of X, NuScenes dataset, urban setting.

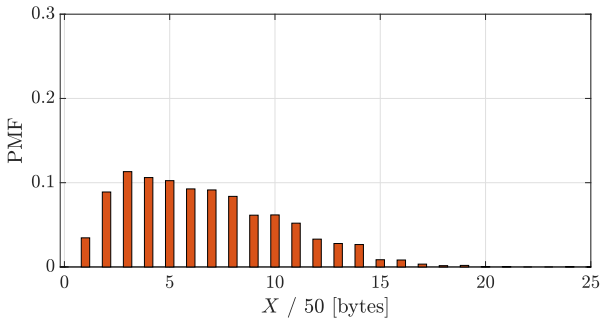

(a) Urban setting.

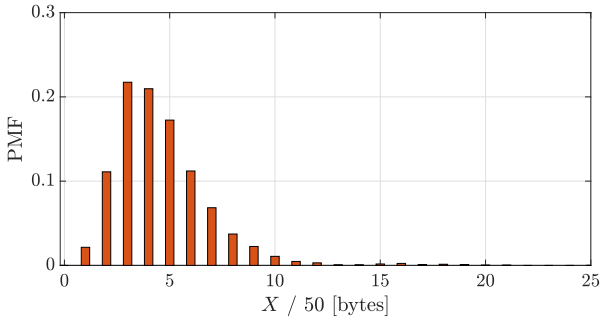

(b) Highway setting.

Fig. 4. PMF of X, Cirrus dataset.

the infrequent occurrences of CPMs with particularly large sizes. Although the data rates in Table I refer to a single CAV, they are useful for estimating the communication burden placed on the communication channel by the future massive deployment of CAVs and cooperative perception applications.

We next focus on the Cirrus dataset, which refers to two scenarios in Palo Alto, California, explicitly termed urban and highway; 55 minutes of continuous recordings are available for the urban environment and 50 minutes for the highway. In this case, we can only consider $T_{genCPM} = 1$ s, as the data sampling rate is 1 Hz. Moreover, owing to the reduced information Cirrus provides, no detection policies can be applied. The comparison between Figs. 4(a) and 4(b) reveals that the PMF of X referring to the urban scenario is less peaked and exhibits a larger variance than the PMF referring to the highway. This is coherent with the previous findings from NuScenes. Notwithstanding that the CAV sensor equipment is different in Cirrus and NuScenes, it is also interesting to confront the PMF of X in Fig. 4(a) against the PMFs previously shown in Figs. 3(a) and 3(b). We conclude that the results obtained from Cirrus fall in between those obtained when the two detection policies (relaxed and restricted) are applied to the NuScenes objects. This corroborates the generality of our analysis and the usefulness of its outcomes.

Hereafter, we leverage the CPM size to determine the average and maximum data rates that characterize a single CAV in the urban and highway scenarios the Cirrus dataset refers to. Albeit the analysis is hampered by the limited number of available data, as well as by the lack of different detection policies and T_{genCPM} settings, the order of magnitude of the values reported in Table II and referring to the urban scenario is comparable to those that were obtained for $T_{genCPM} = 1$ s in Table I.

Given it is primarily N_{obj}, the number of objects included in each CPM which determines the message size, in what follows we concentrate on this variable. For the sake of completeness, Fig. 5 illustrates how N_{obj} affects the CPM size when a different number of onboard sensors is examined. The solid-

TABLE I
AVERAGE AND MAXIMUM DATA RATES, NUSCENES DATASET.

	$T_{genCPM} = 500$ ms		$T_{genCPM} = 1$ s	
	Avg.	**Max.**	**Avg.**	**Max.**
Suburban, Relaxed	4.48 kbit/s	20 kbit/s	3.2 kbit/s	10.4 kbit/s
Suburban, Restricted	3.12 kbit/s	19.2 kbit/s	2.24 kbit/s	9.6 kbit/s
Urban, Relaxed	7.44 kbit/s	42.4 kbit/s	5.36 kbit/s	22.4 kbit/s
Urban, Restricted	4.88 kbit/s	41.6 kbit/s	3.44 kbit/s	21.6 kbit/s

TABLE II
AVERAGE AND MAXIMUM DATA RATES, CIRRUS DATASET.

	$T_{genCPM} = 1$ s	
	Avg.	**Max.**
Urban	2.8 kbit/s	9.6 kbit/s
Highway	2.0 kbit/s	8.8 kbit/s

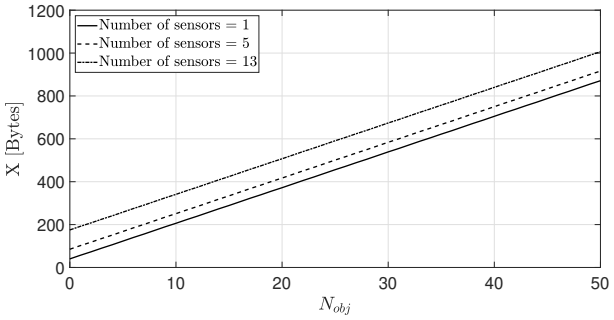

Fig. 5. CPM size as a function of the number of included objects.

line curve refers to the case in which the CAV is equipped with one sensor, the dashed-line curve to the case in which the CAV is equipped with 5 sensors (as in the Cirrus dataset), and the dot dashed-line curve to the case of 13 sensors (as in the NuScenes dataset). It can be noticed that the X dependency on N_{obj} is linear and the CPM size significantly increases when N_{obj} increases.

Next, Figs. 6(a) and 6(b) provide an interesting comparison between the PMF of N_{obj} determined from the Cirrus dataset, and the analogous PMF obtained by simulation in [2], with reference to the highway and urban environments. In detail, the authors of the above study considered the vehicle equipped with a rooftop single camera sensor, with a 360° FoV, and set $T_{genCPM} = 100$ ms. The remarkable difference between the two PMFs is manifest, revealing that: (i) the CAV sensors and

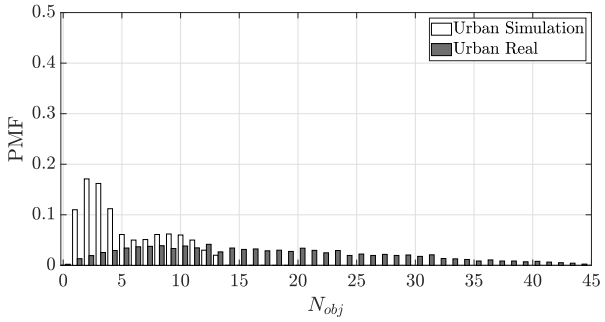

(a) Urban.

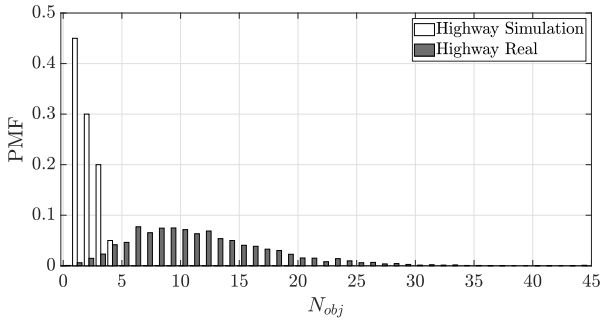

(b) Highway.

Fig. 6. Comparison against the simulative data provided in [2].

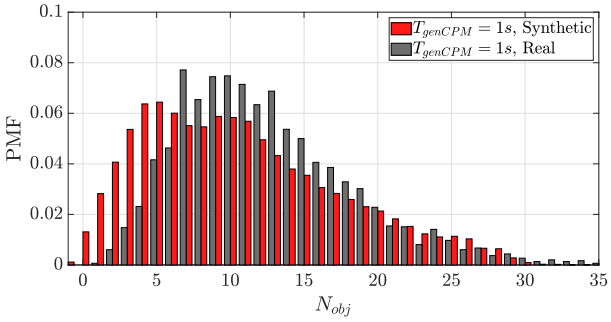

Fig. 7. PMF of the number of objects included in a CPM.

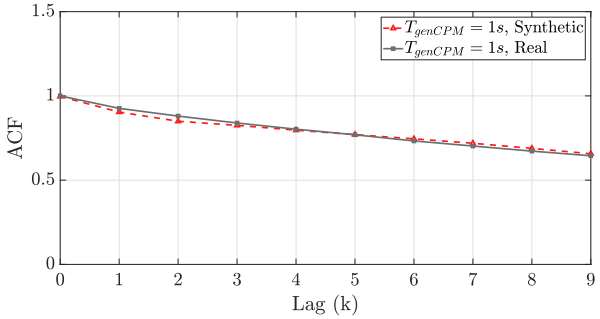

Fig. 8. ACF of the sequences of objects included in a CPM.

sampling capability play a remarkable role in determining the actual viewing perception of the vehicle; (ii) simulation studies may lead to somewhat misleading conclusions, if they do not thoroughly take into account the true CAV equipment.

B. Synthetic Sequences of Objects via TimeGAN

As regards the possibility of generating the artificial sequences of the number of detected objects N_{obj}, we next focus on the exemplary case of the highway scenario and discuss the results obtained by applying the TimeGAN approach to the Cirrus dataset. Fig. 7 reports the PMF of N_{obj} when $T_{genCPM} = 1$ s, as well as the PMF of the synthetic data, corresponding to the optimal setting of the TimeGAN training parameters [11]; namely, the sequence length is 25, the batch size is 60, and the number of iterations is 10^4. Fig. 8 further reports the ACF of the actual sequences of the number of objects included by the CAV in the CPMs, on a 20-second long window, as well as its counterpart for the synthetic sequences.

The obtained KL divergence is 0.0265, the discriminative score is 0.0289 and the predictive score is 0.00521; these are the lowest, i.e., the best values we obtained for the examined range of training parameters. A good match is observed in both figures, allowing us to conclude that the generative model can successfully be employed to obtain temporal sequences of objects, whose number and temporal correlation closely resemble those observed in real recorded sequences.

VI. Conclusions and Future Work

This paper has presented the first study on CPMs generated from real data gathered by AVs traveling in different settings. It has evidenced how datasets conceived for computer vision purposes can be successfully employed to draw accurate information about the data traffic the CAV generates. Through suitable adjustments, the proposed approach can be applied to other annotated datasets to further enrich the analysis.

The investigation of the CPM sequences has revealed that the CPM size is heavily influenced by the level of confidence that characterizes the sensor fusion policy and by the driving scenario, which determines the numerosity of the objects the CAV sees. The examined values of the CPM generation periods were consistent with the actual sampling rate of the scenes and have been shown to impact the message size to a lesser extent. This work has also disclosed the significant distance between previously published data obtained by simulation and those provided by the current realistic analysis.

Lastly, a TimeGAN has been put forth, to model the CAV behavior when transmitting CPM sequences. In the reference highway setting, the generative model satisfyingly replicates the statistical distribution of the number of objects included in the CPMs and captures the temporal correlation between these objects. It is a promising outcome that shows the model is suitable for performing large-scale numerical simulations, where the CPMs simultaneously transmitted by many CAVs have to be generated.

In future works, plans are to extend the current results validating the generative model in different driving environments, also contrasting its behavior to alternative AI-based and statistical methods.

Acknowledgements

The work by Mattia Andreani was funded by the research program “RESearch and innovation on future Telecommunications systems and networks, to make Italy more smART” (RESTART).

At the time of this work, Luca Lusvarghi was with the Department of Engineering “Enzo Ferrari”, University of Modena and Reggio Emilia, Modena, Italy.

References

[1] ETSI TS 103 324 V2.1.1 (2023-06), Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Collective Perception Service; Release 2.

[2] G. Thandavarayan, M. Sepulcre and J. Gozalvez, “Generation of Cooperative Perception Messages for Connected and Automated Vehicles,” in IEEE Transactions on Vehicular Technology, vol. 69, no. 12, pp. 16336-16341, Dec. 2020.

[3] S. Avedisov, A. H. Sakr, T. Higuchi, A. Yamamuro and O. Altintas, “Awareness Assessment of Connected Vehicles in Highway Driving: A Perceived Safety Approach,” in IEEE Vehicular Technology Magazine, vol. 16, no. 3, pp. 129-136, Sept. 2021.

[4] Q. Delooz et al., “Analysis and Evaluation of Information Redundancy Mitigation for V2X Collective Perception,” in IEEE Access, vol. 10, pp. 47076-47093, 2022.

[5] S. Lobo, A. Festag and C. Facchi, “Enhancing the Safety of Vulnerable Road Users: Messaging Protocols for V2X Communication,” 2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall), London, United Kingdom, 2022, pp. 1-7.

[6] nuScenes Full dataset (v1.0), nuScenes by Motional, Mar. 2019. [Online]. Available: https://www.nuscenes.org/nuscenes.

[7] Cirrus Full Dataset, by Volvo, 2020. [Online]. Available: https://developer.volvocars.com/resources/cirrus.

[8] H. Caesar et al., “nuScenes: A Multimodal Dataset for Autonomous Driving,” 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 11618-11628.

[9] Wang Z. et al., “Cirrus: A Long-range Bi-pattern LiDAR Dataset,” 2020. Available: https://arxiv.org/abs/2012.02938.

[10] Vanetza. [Online]. Available: https://github.com/riebl/vanetza.

[11] Y. Yoon, D. Jarrett, M. van der Schaar, “Time-series Generative Adversarial Networks,” in Advances in Neural Information Processing Systems (NEURIPS2019), Vol.32, 2019.