



Towards trustworthy multi-modal motion prediction: Holistic evaluation and interpretability of outputs

Downloaded from: <https://research.chalmers.se>, 2025-12-04 07:19 UTC

Citation for the original published paper (version of record):

Carrasco Limeros, S., Majchrowska, S., Johnander, J. et al (2024). Towards trustworthy multi-modal motion prediction: Holistic evaluation and interpretability of outputs. CAAI Transactions on Intelligence Technology, 9(3): 557-572.
<http://dx.doi.org/10.1049/cit2.12244>

N.B. When citing this work, cite the original published paper.

CAAI Transactions on Intelligence Technology

Special issue Call for Papers

**Be Seen. Be Cited.
Submit your work to a new
IET special issue**

Connect with researchers and
experts in your field and share
knowledge.

Be part of the latest research
trends, faster.

Read more



**The Institution of
Engineering and Technology**

The Institution of
Engineering and Technology

WILEY

ORIGINAL RESEARCH

Towards trustworthy multi-modal motion prediction: Holistic evaluation and interpretability of outputs

Sandra Carrasco Limeros^{1,2,3} | Sylwia Majchrowska^{2,3} | Joakim Johnander^{3,4} |
Christoffer Petersson^{3,5} | Miguel Ángel Sotelo¹ | David Fernández Llorca^{1,6}

¹Computer Engineering Department, Polytechnic School, University of Alcalá, Madrid, Spain

²AI Sweden, Gothenburg, Sweden

³Zenseact AB, Gothenburg, Sweden

⁴Department of Electrical Engineering, Linköping University, Linköping, Sweden

⁵Chalmers University of Technology, Gothenburg, Sweden

⁶European Commission, Joint Research Centre, Seville, Spain

Correspondence

David Fernández Llorca.

Email: david.fernandez-llorca@ec.europa.eu

Funding information

European Commission, Joint Research Center, Grant/Award Number: HUMAINT; Ministerio de Ciencia e Innovación, Grant/Award Number: PID2020-114924RB-I00; Comunidad de Madrid, Grant/Award Number: S2018/EMT-4362 SEGVAUTO 4.0-CM

Abstract

Predicting the motion of other road agents enables autonomous vehicles to perform safe and efficient path planning. This task is very complex, as the behaviour of road agents depends on many factors and the number of possible future trajectories can be considerable (multi-modal). Most prior approaches proposed to address multi-modal motion prediction are based on complex machine learning systems that have limited interpretability. Moreover, the metrics used in current benchmarks do not evaluate all aspects of the problem, such as the diversity and admissibility of the output. The authors aim to advance towards the design of trustworthy motion prediction systems, based on some of the requirements for the design of Trustworthy Artificial Intelligence. The focus is on evaluation criteria, robustness, and interpretability of outputs. First, the evaluation metrics are comprehensively analysed, the main gaps of current benchmarks are identified, and a new holistic evaluation framework is proposed. Then, a method for the assessment of spatial and temporal robustness is introduced by simulating noise in the perception system. To enhance the interpretability of the outputs and generate more balanced results in the proposed evaluation framework, an intent prediction layer that can be attached to multi-modal motion prediction models is proposed. The effectiveness of this approach is assessed through a survey that explores different elements in the visualisation of the multi-modal trajectories and intentions. The proposed approach and findings make a significant contribution to the development of trustworthy motion prediction systems for autonomous vehicles, advancing the field towards greater safety and reliability.

KEYWORDS

autonomous vehicles, evaluation, interpretability, multi-modal motion prediction, robustness, trustworthy AI

1 | INTRODUCTION

The ability of human drivers to predict the motion of other road agents allows us to anticipate potentially dangerous situations and take preventive actions to minimise safety risks. It also allows humans to perform more efficient and comfortable manoeuvres. It is therefore important that autonomous vehicles also have the capability to predict the motion of other road agents, so that they can apply predictive planning approaches and therefore behave in a more human-like manner.

However, predicting future actions and motions of traffic participants is a very complex task, as the behaviour of road agents is influenced by many different variables and interactions [1, 2]. Furthermore, despite the fact that traffic environments are well structured (e.g. street layout, traffic rules), the number of possible future trajectories for each past trajectory for each agent can be considerable, whether for pedestrians, cyclists or vehicles. That is, the problem is multi-modal in nature.

In order to handle this complexity, most of the computational approaches proposed to address multi-modal motion

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *CAAI Transactions on Intelligence Technology* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology and Chongqing University of Technology.

prediction rely on very complex machine learning models which are far from being interpretable. These models are not at human scale and suffer from the characteristic of opacity (i.e., black-box models). Besides, there is no consensus on the most important metrics that should be used to evaluate their performance. Different benchmarks propose different metrics and, in most cases, focus mainly on accuracy, omitting some other relevant aspects such as robustness, diversity, or compliance with traffic rules.

Furthermore, in recent years it has been increasingly accepted that the design of complex learning-based systems must follow certain rules to ensure compliance, not only with traditional safety requirements, but also with general ethical grounds. This approach has recently been referred to as *Trustworthy AI*. It is a concept that encompasses multiple ethical principles, requirements, and criteria to guarantee that AI systems are designed following a human-centred approach and committed to social good [3]. The development of human-centric AI is now a common trend worldwide. For example, at EU level, the High Level Expert Group on AI (AI HLEG) appointed by the European Commission (EC) defined the main horizontal requirements [4] and criteria [5] to develop trustworthy AI systems, including elements such as human oversight, robustness and safety, privacy, and data governance, transparency, fairness, well-being, and accountability. In April 2021, the EC presented the *Proposal for a regulation laying down harmonised rules on AI* (the AI Act [6]) which imposes a set of requirements for AI systems used in high-risk scenarios. Among other things, the AI Act states that the relevant accuracy metrics shall be fit for purpose, and that technical measures shall be put in place to facilitate the interpretation of the outputs of AI systems. At the US level, albeit with a different focus ('algorithms' and 'automatic decision systems' instead of 'AI systems', and 'critical decisions' instead of 'high-risk scenarios'), the *Algorithmic Accountability Act* [7] was introduced in the US Senate and the House of Representatives in February 2022, which also imposes specific requirements on impact assessment, documentation, and performance evaluation of automated critical decision systems.

In this work, we aim to advance towards the design of trustworthy multi-modal motion prediction systems, based on some of the aforementioned requirements. More specifically, we focus on evaluation criteria, including robustness, and interpretability of outputs. The main contributions can be summarised as follows (see Figure 1):

- We comprehensively analyse the evaluation metrics, identifying the main gaps and proposing a new holistic evaluation framework which jointly takes into account accuracy, diversity, and compliance with traffic rules (admissibility).
- We formulate a new method to assess the spatial and temporal robustness of multi-modal motion prediction by simulating noise in the perception system.
- We propose a new intent prediction layer that can be attached to multi-modal trajectory prediction models to enhance the interpretability of the outputs and generate more balanced results in the proposed holistic evaluation framework.

- We assess the interpretability of the outputs and the evaluation framework by means of a survey that explores different elements in the visualisation of the multi-modal trajectories and intentions.

The proposed study is carried out on two of the most widely used datasets in the field (Argoverse [8] and nuScenes [9]), and using two state-of-the-art multi-modal prediction approaches (DenseTNT [10] and PGP [11]).

The remaining of this paper is organised as follows. Section 2.1 presents a literature review on multi-modal motion prediction. In Section 3, we analyse the current evaluation methods for multi-modal motion prediction models and propose a more comprehensive evaluation framework. Section 4 presents the assessment of the robustness of two state-of-the-art multi-modal prediction models in nuScenes and Argoverse. In Section 5 we describe the intention prediction layer to cluster the output predictions into high-level behaviours to enhance the interpretability of the output. Section 6 provides the study on how multi-modality impacts interpretability by means of a survey. Finally, conclusions and next steps are presented in Section 7.

2 | MULTI-MODAL MOTION PREDICTION

2.1 | Methods

Recently, many different approaches were proposed to deal with multi-modal motion prediction [12–17]. The state-of-the-art methods rely on graph models [10, 11, 18–20], since traffic scenarios can be naturally represented as graphs. The graph structure is flexible and efficient because it allows direct and explicit representation of interactions through edges. Usually graph-based models express the features of agents (e.g., vehicles, pedestrians, bicycles) in the nodes, and the relationships between agents as edges [20]. There are several architectural design decisions that must be made in order to effectively represent the input data, model the interaction, and finally represent the output trajectory distributions.

2.1.1 | Representation of high-definition maps

Methods for motion prediction need to effectively represent both geometric information (static scene elements) and traffic agents (dynamic scene elements). The standard raster representation encodes the world as a stack of bird's eye view (BEV) images, also called high-definition (HD) maps [12, 13, 21, 22]. This approach is straightforward as all the different types of input information (e.g., road configuration, state history of agents, and spatial relationships) are unified in a multi-channel image, allowing the use of a Convolutional Neural Network (CNN). However, this approach is limited by narrow receptive field of standard CNNs. Hence, rasterised representations have difficulties in modelling long-distance interactions.

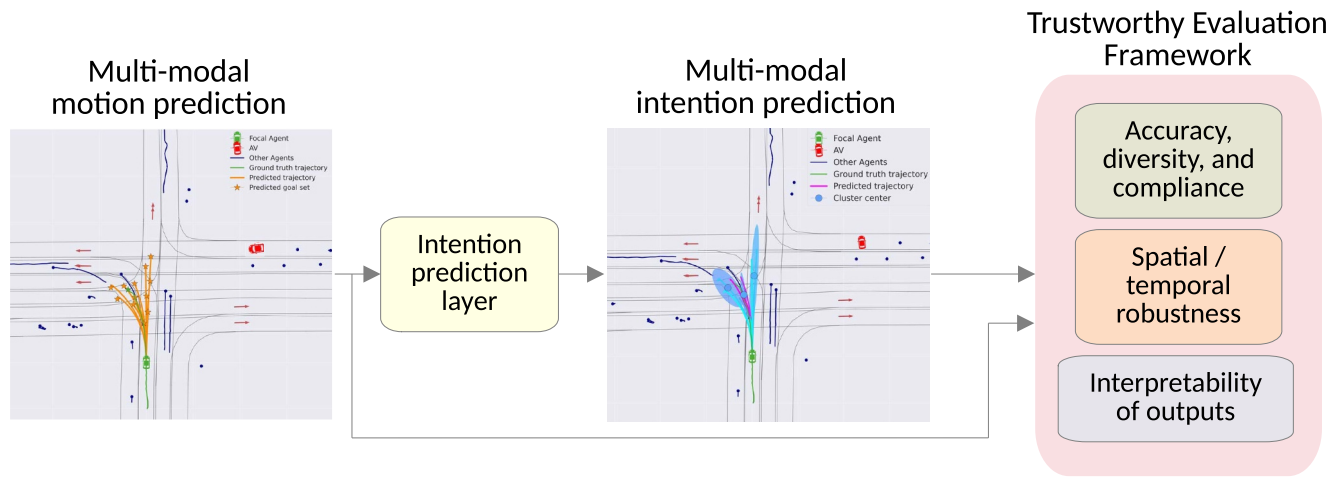


FIGURE 1 Towards trustworthy multi-modal motion prediction. The intention prediction layer enhances the interpretability of the results. The proposed evaluation framework includes a holistic set of metrics and a robustness analysis.

Alternatively, less expensive graph-based representation can be used in the form of polylines (e.g., lanes, crosswalks, and boundaries), which represent piecewise linear segments [10, 16, 18, 19, 23, 24]. In this representation, long-range dependencies are effectively and efficiently modelled. For example, VectorNet [10, 18, 19] represents the lane segments and agents histories as vectors connected end-to-end. As an improvement, the Lane-based Trajectory Prediction (LTP) method [23] encodes the map through connected sliced lane segments to more precisely model the intention. Other approaches [11, 25–27] construct an additional lane graph, leveraging temporal and spatial correspondence of the vectorised scene context.

2.1.2 | Interaction modelling

There are two types of interactions that need to be modelled. First, the encoding of temporal sequential data, which is typically accomplished via Recurrent Neural Networks (RNNs)—such as the Gate Recurrent Unit (GRU) [28] or the Long Short-Term Memory (LSTM) [29]—or temporal convolutions. Second, the interactions between the relevant agents and the environment is modelled via an attention mechanism [30]. Interaction modelling is closely related to the method used for scene encoding.

Based on the polyline representation, VectorNet [10, 18, 19] utilises only self-attention modules to directly learn the interactions between all the sub-graphs in the environment. Further, goal-oriented lane attention [31] emphasises the relationship between agents and lanes. The split-joint attention mechanism, used by LaneGCN [26], captures the complex topology of lane graphs and long range dependencies. Deo et al. [11] propose a Prediction via Graph-based Policy (PGP) model, where interactions are modelled via lane-graph traversals. This approach combines discrete policy roll-outs with a lane-graph subset decoder, conditioning each prediction on the driver's goals. Finally, the network predicts trajectories by selectively attending to node encodings along paths traversed by the policy and a sampled latent variable.

2.1.3 | Multi-modal output

To account for environmental uncertainty, models predicting future trajectories should represent multi-modal paths [21, 31, 32]. This can be represented by implicitly modelling multi-modality as latent variables or explicitly proposing to generate multiple trajectory proposals. The first approach uses Gaussian Mixture Models (GMMs) [33] or Mixture Density Networks (MDNs) [34] to generate distributions for possible trajectories. In next step, it uses Variational Autoencoders (VAEs), recently also their modified version as Conditional Variational Autoencoder (CVAEs) [35], or GANs [33, 36, 37] to sample various future modes from latent variables. The main drawback of this approach is that the obtained predictions cannot be interpreted unambiguously, which lowers the understanding of the model's predictions.

On the other hand, the proposal-based methods design various possible proposals and thus separately solve the tasks of intention prediction and motion prediction. Some methods [10, 19] sample proposal points around the lane centreline to capture detailed information. The proposed sampling is based on carefully developed rules, the restriction on sharing proposals between different agents is one of them. In contrast, [23] uses lane segments from the map as proposals that can clearly describe the fine-grained intentions of agents and be shared globally.

From predicted paths, the most reliable trajectories are selected. Here, a variation of Non-maximum Suppression (NMS) algorithm is widely used. The traditional version of this method is employed in TNT [19], in which each trajectory is sorted according to the method scoring process. Then, the distance between different trajectories is calculated and a diverse set of trajectories with high scores is selected. Similar to TNT, LaneRCNN [27] treats a lane segment as an anchor and output each anchor's probability, using NMS to remove too close duplicate goals. The solution has its drawbacks, as a fixed threshold does not allow the model to maintain the balance between accuracy and multi-modality of output. In contrast,

DenseTNT [10] is an anchor-free goal-based model, which generates a set of predicted goals without relying on the heuristic anchors.

2.2 | Datasets

Experiments presented in subsequent sections are conducted on two publicly available trajectory forecasting benchmarks: The Argoverse v1.1 Motion Forecasting Dataset [8] and nuScenes [9]. These two datasets are the most widely used in the motion prediction task. The datasets differ from each other in many parameters, for example, number of scenes, total driving time, number of objects, interaction complexity, or prediction time horizon. The prediction time horizon, in particular, plays a crucial role. Long-term forecasting is inherently more difficult, and multi-modal forecasting appears to be more successful in this case. Correctly determining the driver's intentions and the future trajectory of the vehicle is essential for safe planning. It strongly depends on the overall quality of the dataset on which the algorithms were trained, so it is important to select appropriate data for future applications.

Argoverse [8], released in 2019 by Argo AI, is composed of 323 557 real-world driving sequences. Driving scenarios were collected in two American cities: Miami and Pittsburgh. Each example consists of 2 s of historical state and 3 s of future state, which is sampled with a frequency of 10 Hz. The entire collection of scenarios totals 320h. In each scene, one actor of interest is specified whose future movement is to be predicted. For training and validation, focal agent history, location histories of nearby (social) actors, together with HD map features are also provided. The semantic HD map provided is simplistic and consists of lane-based polylines.

nuScenes [9], released in 2020 by Motional, is a large-scale data collection for multi-agent trajectory forecasting with 1000 scenes recorded in Boston and Singapore, where right-hand and left-hand traffic rules apply respectively. Each scene is annotated at a frequency of 2 Hz and is 20 s long. It contains up to 23 semantic object classes, including vehicles, bicycles, and pedestrians as possible tracked agents, as well as HD semantic maps with 11 annotated layers. For this dataset, each agent has 2 s of observed trajectory and the prediction horizon is set to 6 s.

Overall, Argoverse—which tends to represent short, independent episodes—is more than 50 times larger in terms of total time than nuScenes. Despite the impressive size of the Argoverse dataset, previous works have observed that the recorded trajectories mostly represent straight-line trajectories within a full 5 s window [24, 38]. This makes the dataset much less challenging and diverse than nuScenes. The datasets differ also in the schematic representation of the HD maps. Argoverse presents a much simpler representation, while nuScenes is more comprehensive but complex. We simplified nuScenes representations to focus on the most representative information for the conducted study.

3 | EVALUATION FRAMEWORK

In this section, we revise the evaluation procedures of multi-modal prediction approaches. The presented metrics are differentiated between measures of precision, diversity, and admissibility (see Table 1). Diversity refers to the degree of coverage of the output distribution. We seek to have diversity across distribution modes rather than having multiple candidates representing a single intention (mode). It is important to evaluate, however, the admissibility of such predictions. They need to comply with the traffic rules.

3.1 | Accuracy evaluation

- **Best-of-K Average Displacement Error (minADE):** The minimum point-wise L2 distance to the ground-truth trajectory over all predicted trajectories.
- **Scene minADE:** Used when considering the prediction of multiple agents simultaneously. It measures the joint minimum L2 distance between predictions and ground-truth over all the agents in the scene.
- **Best-of-K Final Displacement Error (minFDE):** Lowest L2 distance between the K predicted endpoints and the ground truth endpoint at the prediction horizon over all predicted trajectories.
- **Scene minFDE:** To some extent, similar to scene minADE. It measures the joint minimum FDE over the whole scene.
- **Best-of-K Miss Rate (MR):** Percentage of predictions whose maximum pointwise L2 distance between the prediction and ground truth is greater than a threshold. For Argoverse and nuScenes benchmarks a threshold of 2.0 m is considered. However, they define this metric in different ways. Whereas the nuScenes benchmark uses the above definition, Argoverse computes MR as the number of scenarios where none of the forecasted trajectories are within 2 m of the ground-truth endpoint. Waymo [39] differentiates between lateral and longitudinal thresholds, which scale depends on future time horizon and initial speed.
- **Heading error:** The difference in the heading angle between the predicted trajectory and the ground truth trajectory at the final point.
- **Mean Average Precision (mAP) [39]:** The area under the precision-recall curve by applying confidence score thresholds across a validation set. They use the same definition of a miss as defined for MR, considering any missed prediction as false positive. Only one true positive is allowed for each prediction—assigned to the highest confidence prediction. They report the final mAP averaged over eight different semantic buckets or driving behaviours, that is, straight, straight-left, straight-right, left, right, left u-turn, right u-turn, and stationary. This makes the evaluation more balanced, since some of these trajectory shapes are much more infrequent than others.
- **Soft Mean Average Precision (Soft mAP) [39]:** Similar to mAP, but additional matching predictions, other than the highest confidence one, are not penalised.

TABLE 1 Reviewed metrics categorised in terms of accuracy, diversity, and admissibility. p corresponds to the probability of the best forecasted trajectory. Symbol - refers to *no units*, # refers to *number of*.

Metric		Unit	Description
Accuracy	Non-probabilistic	m (4)	The average of pointwise L2 distances between the predicted trajectory and ground truth over the k most likely predictions.
		m (4)	The L2 distance between the final points of the prediction and ground truth, over the k most likely predictions.
		% (4)	Percentage of scenarios where none of the forecasted trajectories for an agent are within 2.0 m of ground truth according to endpoint error.
		rad (4)	Difference in the heading angle between the predicted trajectory and the ground truth trajectory at the final point.
		- (f)	Area under the precision-recall curve by applying confidence score thresholds across a validation set. True and false positives are defined using the definition of MR. Only one true positive is allowed for each agent, assigned to the highest confidence prediction.
Scene-level		- (f)	Similar to mAP, where additional matching predictions, other than the highest confidence one, are not penalised.
		m (4)	Joint minimum L2 distance between predictions and ground-truth over all the agents in the scene.
		m (4)	Joint minimum L2 distance between the K predicted endpoints and the ground truth endpoint at the prediction horizon over all predicted trajectories.
		m (4)	Similar to minADE, adding $\min(-\log(p), -\log(0.05))$ to the average L2 distance.
		m (4)	Similar to minFDE, adding $\min(-\log(p), -\log(0.05))$ to the endpoint L2 distance.
Probabilistic		% (4)	Similar to Miss Rate, taking $(1.0 - p)$ as the contribution, instead of 0.0, when the endpoint error for the best forecasted trajectory is less than 2.0 m.
		m (4)	Similar to minFDE, adding $(1.0 - p)^2$ to the endpoint L2 distance.
		m (4)	Similar to minADE, adding $(1.0 - p)^2$ to the average L2 distance.
		- (4)	The average negative log likelihood of the ground truth trajectory as determined by a kernel density estimate over output samples at the same prediction timestep.
		m/m (f)	Ratio avgFDE to minFDE as a measure of the spread of the predictions in Euclidean distance.
Diversity		# (f)	Average number of different lanes reached.
		rad ² (f)	Variance of the final heading for the different output modes.
		m (4)	Minimum value among the L2 distance between centrelines of possible lane candidates and each predicted mode.
		m/s ² (f)	Variance of the average speeds.
		m/s ² (f)	Variance of the average accelerations.
Admissibility		% (4)	Fraction of trajectories that are outside the drivable area.
		# (f)	Proportion of pixels occupied by the predicted trajectories within the drivable area.
			(Continues)

TABLE 1 (Continued)

	Metric	Unit	Description
Social interactions	Drivable Area Compliance (DAC)	# (↑)	If a model produces k possible future trajectories and m of that exit the drivable area at some point, the DAC for that model would be $(n - m)/n$.
	Scene consistency rate (SCR)	% (↓)	Percentage of predicted samples that overlap in the scene looking at the IOU between the future BEV-defined bounding boxes of each pair of agents.
	Overlap Rate	% (↓)	Number of overlaps divided by the total number of agents. A single overlap is defined if any of these trajectories overlaps at any time with any other agent at the prediction time step.
Traffic rules	Oncoming Traffic Direction (OTD)	% (↓)	Ratio of modes going in the oncoming traffic direction.

Multi-modal prediction models are usually benchmarked using Best-of-K metrics. These metrics, although useful for deterministic regressors, only take into account a single output of an arbitrary number K , representing a limited part of the model output distribution. Hence, they are not able to compare the distributions produced by multi-modal models, neglecting the assessment of variance and multi-modality.

Furthermore, these metrics only assess the quality of the underlying marginal distribution per agent. Best-of-K ADE takes the trajectory sample that is closest to the ground truth of each agent in an independent manner. In this way, we are not measuring scene consistency in the predictions. It would be possible to have a low minADE by predicting high-entropy distributions that are not consistent at the scene level. In [40], the authors propose scene-level sample metrics to assess how well the output modes capture the joint distribution over future motions.

3.2 | Probabilistic evaluation

In order to make a fair assessment of the probabilistic capabilities of a model, we need to measure its ability to capture the underlying uncertainty distribution of future motions. In the following, we list some of the metrics used in the literature [8] to evaluate probabilistic prediction models. p being the probability of the best forecasted trajectory:

- Probabilistic minimum Final Displacement Error (p-min-FDE): Similar to minFDE. It adds $\min(-\log p, -\log 0.05)$ to the endpoint L2 distance.
- Probabilistic minimum Average Displacement Error (p-minADE): Similar to minADE. It adds $\min(-\log p, -\log 0.05)$ to the average L2 distance.
- Probabilistic Miss Rate (p-MR): Similar to Miss Rate. It takes $1.0 - p$ as the contribution, instead of 0.0, when the L2 distance between the maximum pointwise L2 distance between prediction and ground truth is greater than the threshold.
- Brier minimum Final Displacement Error (brier-minFDE) [41]: This metric is similar to minFDE, but we add $(1.0 - p)^2$ to the endpoint L2 distance.
- Brier minimum Average Displacement Error (brier-min-ADE) [41]: This metric is similar to minADE, but we add $(1.0 - p)^2$ to the average L2 distance.
- Negative Log Likelihood (NLL) [42]: The average negative log likelihood of the ground truth trajectory as determined by a kernel density estimate over output samples at the same prediction timestep.

It is important to note that different benchmarks consider a different number of modes to evaluate prediction accuracy. For instance, Argoverse and Waymo consider the 6 most likely trajectories, whereas nuScenes takes 5 and 10 modes. For probability-based metrics in Argoverse, they take the 6 most likely trajectories and normalise the probabilities before computing the metrics.

There is also a difference in the prediction horizon considered. Both Argoverse and nuScenes provide 2s of trajectory history. However, the Argoverse prediction time horizon is 3s, whereas nuScenes considers 6s. Waymo [39] provides tracks for the past 1 s and considers three different evaluation times: 3, 5, and 8 s into the future.

In order to sort their leaderboards, there is no consensus on the most important metric. nuScenes ranks the different approaches according to minADE₅ ($K = 5$). Argoverse sorts them according to brier-minFDE at $K = 6$, assuming a uniform distribution for approaches that do not provide a probability associated to the predictions. In previous competitions, they used MR and minFDE at $K = 6$. Waymo leaderboard ranking is built based on the average Soft mAP across evaluation times, while MR is used as a secondary metric.

3.3 | Diversity evaluation

Evaluating the diversity of multiple predictions is required to properly assess the multi-modal predictive capabilities. However, this evaluation is far less explored in the literature, resulting in models that, although obtaining good results in terms of accuracy, suffer from mode collapse and do not show real multi-modality. In the following, we list metrics used to evaluate the diversity of the predictions:

- Lateral diversity metrics: The authors of [11] report the average number of different lanes reached as a measure of lateral diversity, as well as the variance of the final heading for the different output modes. [38] defines *minLaneFDE*, which captures both the quantity and quality of diversity of multiple outputs based on the centrelines of reference lanes. It computes the minimum value among the L2 distance between centrelines of possible lane candidates and each predicted mode.
- Longitudinal diversity metrics: Deo *et al* [11] reports the variance of the average speeds and accelerations for the different output modes.
- Ratio of avgFDE to minFDE (RF): Proposed in [43], it measures the spread of the predictions in Euclidean distance.

A large average L2 error implies that the predictions are spread out, while a small minimum L2 error implies that at least one of the predictions has high precision. The authors in [43] follow this intuition and propose to compute the ratio of avgFDE to minFDE to capture diversity in the output. However, this metric can fail to differ between longitudinal and lateral diversity in some scenarios. In addition, this metric may be better for models that are worse in terms of accuracy and admissibility when most modes are failing modes. A straightforward way to assess lateral diversity can be to measure the variance of the final heading for the K possible outputs. When having access to the lane information, the average number of final lanes reached is a good measure of mode diversity. On the

other hand, longitudinal diversity can be measured by the variance of the final speed and acceleration for the K outputs.

3.4 | Admissibility evaluation

Finally, assessing the compliance of the predictions with the driving scene is essential to evaluate the quality of the output and to ensure a safe motion planning in complex driving scenarios. In the following, we list the metrics used to evaluate the admissibility of the predictions:

- Off-road rate: Computes the fraction of trajectories that are off-road, outside the drivable area.
- Drivable Area Occupancy (DAO) [43]: Measures the proportion of pixels occupied by the predicted trajectories within the drivable area. This metric can be used together with RF to get a more compliant measure of diversity.
- Drivable Area Compliance (DAC) [8]: Measures extreme off-road predictions that are not admissible. If a model produces K modes for future trajectories and M of those going off the drivable area at any point, the DAC for that model would be $(N - M)/N$.
- Scene consistency rate (SCR) [40]: Measures the percentage of predicted samples that collide (overlap) in the scene in order to evaluate social-consistency. A collision is detected by comparing the IOU between the future BEV-defined bounding boxes of each pair of agents in the scene with a small IOU threshold.
- Overlap Rate [39]: Similar to the previous metric, it takes the highest confidence prediction from each agent and computes the total number of overlaps divided by the total number of agents. A single overlap is defined if any of these trajectories overlaps at any time with any other agent at the prediction time step.

Off-road rate, DAO, and DAC are useful for measuring the proportion of modes that go off-road. The last two metrics evaluate the ability of the model to capture social interactions. However, to the best of our knowledge, there are no metrics that assess compliance with traffic rules. To this end, we propose to measure the ratio of modes going in the oncoming traffic direction $OTD = \frac{\text{count}(\text{traj}_{OTD})}{K}$.

Autonomous driving is a high-stake and safety-critical application. As such, it is of utmost importance to thoroughly evaluate each of its intermediate systems, including the trajectory prediction stage that will be an essential input for safe and efficient planning. A comprehensive evaluation and interpretation of the performance of the prediction model is needed in terms not only of precision, but also diversity and admissibility. Simultaneous examination of all these dimensions provides a holistic evaluation framework for the assessment of multi-modal motion prediction. In practice, the choice of which metrics to evaluate will depend on the particular needs of the system and the specific application context. Our framework serves as a guide to help the user make an informed

decision about which metrics are most relevant to their needs and to provide a basis for comparison with other models.

4 | ROBUSTNESS ANALYSIS

The motion prediction task usually assumes perfect perception, that is, that the ground-truth past trajectory of all actors is given and that the map information is available. However, in practice, self-driving vehicle perception systems have noise that will translate into false negatives, false positives, and incomplete agent tracking or id switches. We therefore conduct a robustness study, in which perception system noise is simulated.

We perform an ablation for PGP [11] as having an object detector with 80%, 90%, and 95% recall—the latter being a realistic setup under good conditions. This ablation is done in an independent manner for dynamic objects detection and lane detection.

Each agent or lane segment is masked following a Bernoulli distribution with a probability of 20%, 10% and 5% respectively. Table 2 shows the results for both lanes and dynamic agents in terms of $minADE_k$, $miss_rate_k$ —metrics used for ranking nuScenes benchmark—and Behavioural Cloning (BC).¹

We found the PGP model to be quite robust to failures in the detection of dynamic agents. However, when we introduce noise in the perception of lanes, performance is drastically decreased, specially in the BC metric, given the importance of the lanes for this metric. This is probably due to the fact that this model largely exploits lane information as a strong inductive bias, which contains both the direction of traffic flow and legal paths for each agent.

We perform a second experiment where we analyse the effect of not detecting some frames of the agents that interact with the focal agent. Results are shown in the second last row. We consider an interaction if their trajectory intersects with a 20 m radius of the target agent. Each frame has a 50% probability of being detected. Results do not show a noticeable drop in performance in this case either. Finally, we perform a final experiment where we simulate that no dynamic agent is detected. Results are shown in the last row. Given these results, we can state that the model relies heavily on lane information as well as on the focal agent's past trajectory.

In order to evaluate if these results generalise to a different scenario and model, we perform the same experiments for DenseTNT in the Argoverse dataset with $K = 12$. Table 3 shows the results of the robustness analysis for DenseTNT. In the same way as in the previous scenario, failing to perceive the lanes has a more detrimental effect than failing to perceive the dynamic agents in terms of accuracy. Again, the masking of all agents in the scene decreases the performance, but it still maintains a reasonable level. Noise in the perception of lanes reduces diversity, while masking of dynamic agents seems to increase diversity—see second-last row in Table 3. The reason for this behaviour is that the fewer lanes perceived, the fewer

plausible trajectories. Not perceiving certain agents, however, could cause the AV to go to an occupied space, increasing the diversity in the prediction at the cost of raising the risk of a collision. Admissibility decreases when lanes are masked for the same reason, while not perceiving agents in the scene has no effect on this aspect, as expected.

The behaviour observed in this analysis can have several explanations. First, the historical information of the focal agent in terms of position, velocity and acceleration is already very informative when inferring the motion of the vehicle. Second, most scenarios lack interactions that drastically impact the future trajectory of the focal agent. Even though nuScenes is one of the most complex trajectory forecasting datasets and it includes several highly interactive scenarios, this is still insufficient. This is one of the main weaknesses of current motion prediction benchmarks. In the most safety critical situation, other road users should play a crucial role. However, this is not reflected in the current benchmarks, which should cover a wide range of edge cases. These rare occurrences are easily missed and thus are often missing in datasets. Humans are naturally proficient at dealing with these extreme cases, but this is not true for autonomous systems. Therefore, we need to deal with it carefully. Another conclusion we can draw from this analysis is the importance of dynamic object detection. Object detectors are usually trained on single images. However, this can be sub-optimal. Humans exhibit an understanding of the dynamics of scenes. If objects move against a static background, we can detect them quite easily, despite darkness, rain or other occlusions. A dynamic object detection system would prevent missing objects in specific frames, thus avoiding dragging or carrying this error into the next stages of the pipeline.

5 | INTENTION PREDICTION

Understanding the intention of the surrounding road agents is most relevant to mid- and long-term prediction and decision making. In order to drive through dynamically changing traffic scenarios, a multi-modal intention prediction module is necessary to adapt to the different scenarios. Intention prediction differs from motion or trajectory prediction in that it corresponds to discrete high-level behaviours, semantically different from each other, which we can consider modes of the future trajectory distribution. Many trajectory prediction models are not inherently multi-modal and suffer from mode collapse. We wish to disambiguate the output and disentangle these modes into clear high-level intentions. In addition, this may contribute towards the interpretability of the overall system since it is more aligned to how humans think while driving.

Another potentially important advantage of intention prediction is that it is much less sensitive to the actions of the AV, compared to the more detailed task of trajectory prediction in which the precise future trajectory of the agent will depend on the actions of the AV if they are ‘interacting’. This implies a greater potential for the intention prediction task to be dealt with in an open-loop manner.

¹Robustness Analysis code repository.

TABLE 2 Robustness analysis of PGP in NuScenes dataset. $K = 10$ for all experiments.

Object	Experiment	minADE ₅	minADE ₁₀	Miss rate ₅	Miss rate ₁₀	BC
Lanes	Baseline	1.30	0.97	0.52	0.34	1.91
	Recall 100%, $K = 10$					
	Recall 95%	1.36	1.01	0.55	0.38	2.51
	Recall 90%	1.41	1.04	0.57	0.40	3.0
	Recall 80%	1.54	1.12	0.60	0.43	4.01
Dynamic agents	Recall 95%	1.30	0.97	0.52	0.34	1.93
	Recall 90%	1.31	0.97	0.53	0.35	1.93
	Recall 80%	1.32	0.97	0.53	0.34	1.94
	Recall 50% temporal	1.34	0.98	0.54	0.36	1.96
	No detected agents	1.42	1.01	0.58	0.39	2.07

Note: Values in bold represent the best performance results within the comparative analysis.

TABLE 3 Robustness analysis of DenseTNT in Argoverse dataset. $K = 12$ for all experiments.

Object	Experiment	minADE	minFDE	Miss rate	AvgFDE	RF	DAC
Lanes	Baseline	0.763	1.078	0.043	5.447	5.053	0.971
	Recall 100%, $K = 12$						
	Recall 95%	0.777	1.110	0.049	5.499	4.954	0.971
	Recall 90%	0.799	1.151	0.059	5.577	4.849	0.971
	Recall 80%	0.858	1.281	0.084	5.823	4.546	0.969
Dynamic agents	Recall 95%	0.763	1.081	0.043	5.459	5.049	0.971
	Recall 90%	0.764	1.082	0.044	5.482	5.066	0.971
	Recall 80%	0.765	1.083	0.044	5.522	5.099	0.971
	Recall 70%	0.768	1.091	0.045	5.571	5.106	0.971
	Recall 50% temporal	0.773	1.107	0.046	5.933	5.359	0.971
	No detected agents	0.803	1.151	0.057	6.113	5.311	0.971

Note: Values in bold represent the best performance results within the comparative analysis.

Most studies in the field of motion prediction work on trajectory prediction and only a few on intention prediction [44, 45], which frame the task as a classification problem. However, these rely on predefined trajectories obtained by hand-crafted principles, failing to capture comprehensive representations for the future distribution. These methods lack a consistent evaluation, which makes them suboptimal and lag behind state-of-the-art regression and generative models. In this section, we explore a new formulation for the intention prediction framework, by using a simple post-hoc approach that could be added on top of different state-of-the-art multi-modal motion prediction methods.

We extend the DenseTNT model [10] to instead perform intent prediction.²

We observed that the set of K potential output goals for the focal agent are not inherently different. Oftentimes, they are right next to each other or almost overlapping. The reason for this is most likely that there is uncertainty in the motion

profile. It would be more convenient and intuitive if each of the K locations were different, disentangling the longitudinal uncertainty from the intention. Another problem we detected in this model is that often the predicted goals are not admissible or rule abiding, falling off the road or in lanes going in the opposite direction. It would also be desirable to have a probability associated with each mode, to facilitate the subsequent decision making.

A simple approach to solve this would be to cluster the different outputs of DenseTNT into intentions. As a first step, we trained DenseTNT to output 12 goals instead of 6—by default—optimising the miss rate metric instead of FDE, since we believed these two changes would lead to better coverage of the output distribution. Indeed, it showed higher diversity and a reduction of failure modes.

The model outputs a set of goals $\mathcal{G} = \{g_1, \dots, g_N\}$, where each goal $g_n \in \mathbb{R}^2$ is a point in birds-eye-view. In order to capture different intentions, we cluster the goals, forming a set of clusters $\mathcal{C} = \{c_1, \dots, c_K\}$. Each cluster, c_k , contains three components,

²DenseTNT-Intent code repository.

$$c_k = (\gamma_k, \mu_k, \Sigma_k), \quad (1)$$

where $\gamma_k \in [0, 1]$, with $\sum_k \gamma_k = 1$, is the probability of a goal ending up in cluster k . $\mu_k \in \mathbb{R}^2$ is the mean position of the cluster and $\Sigma_k \in \mathbb{R}^{2 \times 2}$ is the covariance of the cluster. Note that we expect the covariance to be high along the lane, corresponding to the variability in speed between agents, and low orthogonally to the lane. We propose a straightforward and intuitive approach to find \mathcal{C} .

5.1 | Cluster creation

Clustering is an NP-hard problem and commonly used algorithms, such as K-means clustering or expectation maximisation, are computationally expensive and in some cases prone to degenerate solutions. In our setting, however, we have a good partitioning of the bird's-eye-view plane available already: namely the lane segments in the HD-map. We therefore propose to create clusters based on the lane segments. In the following, we describe the heuristic used to compute the clusters. For each predicted goal g_n :

- 1) Obtain all lane segments $L_n = \{l_1, \dots, l_M\}$ that intersect with a radius $r = 20\text{m}$ from the agent based on the Manhattan distance. The bounding boxes of small point clouds (lane centreline waypoints) are precomputed in the map. If no lanes are found, we double the search radius r .
- 2) Estimate the confidence δ of each lane based on the distance to its closest waypoint, p_{nm} . Keep those lanes whose closest waypoint lies within a threshold radius t of 2.5 m.

$$\delta_{nm} = 1 - \frac{\min(\|g_n - p_m\|_2)}{t} \quad (2)$$

- 3) Compute the angle between the agent's heading for mode n and the lane direction, a_{nm} . Discard those lanes whose direction differs from the agent's heading by more than 45° . This ensures that we cluster only those goals that follow the current lane direction.

Formally, let α_{nk} correspond to the probability that the goal g_n belongs to cluster c_k . We intentionally let the assignment be soft. This makes it possible for a single goal to be assigned to multiple clusters, which is sensible under uncertainty. Taking the retrieved lanes L_n for each goal g_n , we compute α_{nk} as follows:

- 1) If two lanes in L_n do not merge or one is the successor of the other, then they belong to different clusters.
- 2) The probability α_{nk} , where lane l_m belongs to cluster c_k , is computed based on the distance to the closest waypoint p_{nm} and the angle a_{nm} .

$$p_{nk} = \delta_{nm} \frac{1}{a_{nm}}, \quad (3)$$

$$\alpha_{nk} = \frac{e^{p_{nk}}}{\sum_k e^{p_{nk}}}.$$

This heuristic is computationally efficient and guarantees that different intentions—such as staying on lane versus cutting out of the lane—are represented by different clusters. We repeat the process for each goal g_n , grouping the lanes that belong to the same cluster.

Then, we compute

$$\mu_k = \frac{\sum_n \alpha_{nk} g_n}{\sum_n \alpha_{nk}}, \quad (4)$$

$$\Sigma_k = \frac{\sum_n \alpha_{nk} (g_n - \mu_k)(g_n - \mu_k)^T}{\sum_n \alpha_{nk}}. \quad (5)$$

In order to compute the cluster probability, γ_k , we exploit the heatmap produced by DenseTNT. Taking the score of goal g_n , s_n :

$$p_k = \sum_n s_n \alpha_{nk},$$

$$\gamma_k = \frac{e^{p_k}}{\sum_k e^{p_k}}. \quad (6)$$

5.2 | Cluster visualisation

For the survey study, we visualise the clusters by considering a *hard* clustering, assigning each goal to its most probable cluster.

In Figure 2, we showcase different scenarios. The AV is represented in red and the focal agent, whose trajectory we want to predict, is represented in green. Other agents are represented with blue dots. The ground truth trajectory is shown in green.

In the first row, the outputs of the DenseTNT model are visualised in orange. The final predicted goals are represented by orange stars. The arrows show the direction of the lanes.

The clustered output is shown in the second row. The probabilities of the output trajectories are mapped to the coloured bar shown on the right. The averages of the intention clusters, whose probability follow the same mapping, are represented by coloured circles. The uncertainty in the motion profile for each intention is represented by the confidence ellipse of its covariance, visualised as a shaded contour. As expected, the covariance is high along the lane direction and low orthogonally, showing higher uncertainty in velocity. To deal with cases where only one goal is assigned to a cluster, we apply a fixed covariance based on our previous knowledge of 2 m along the lane and 0.05 m orthogonal to the lane (left-most lane cluster in Figure 2).

In the first column, three intentions are detected, one for each lane, being the second the most probable one. In the second column, the model outputs two intentions: follow and change lane. The goal that go off the road is discarded and not considered for the clustering. In the third case, we can observe how this method eliminates spurious goals that do not follow traffic rules, going to lanes that go in the opposite

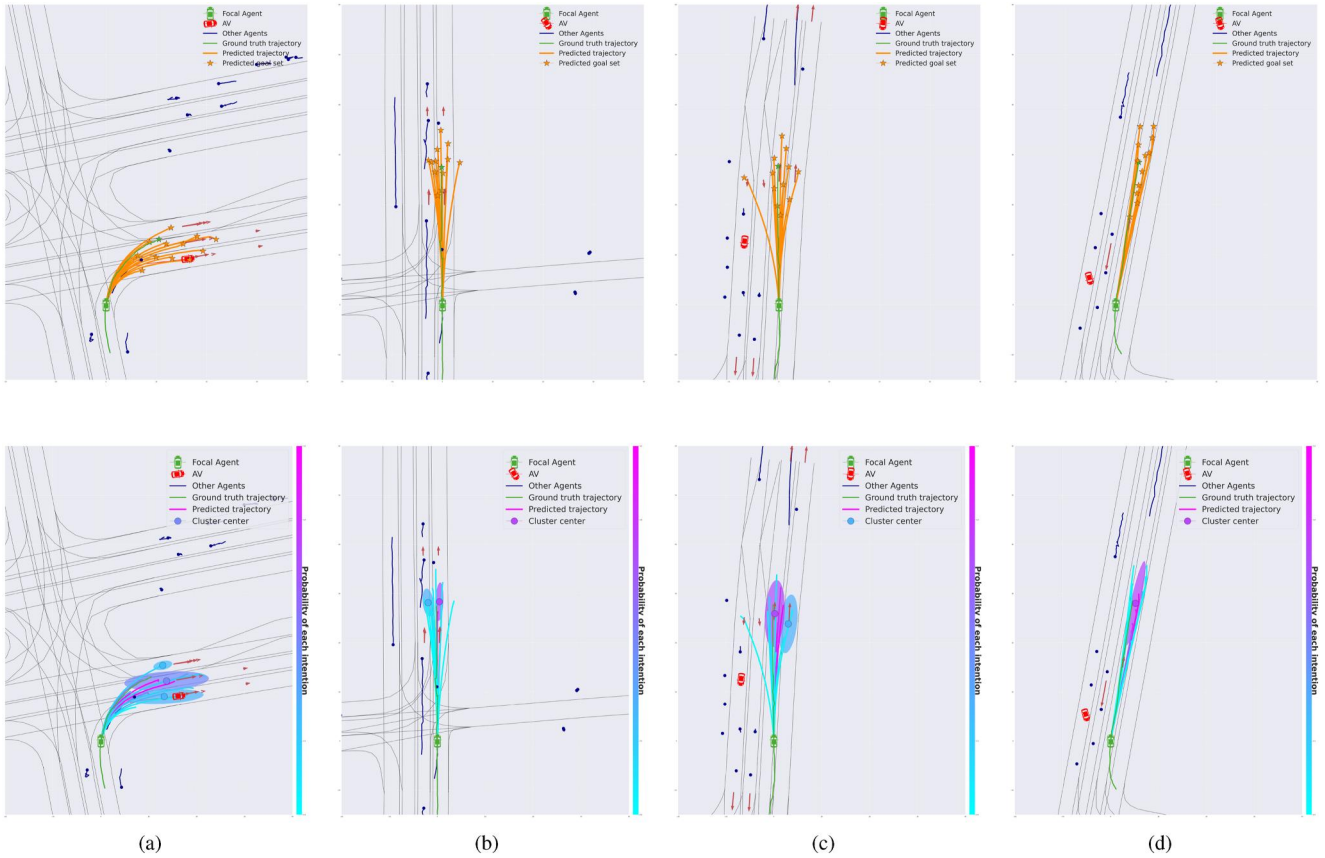


FIGURE 2 Qualitative evaluation: DenseTNT-Intent (*second row*) clusters the output of DenseTNT (*first row*) into high-level intentions with a probability associated with each cluster following the colour bar on the right. The circles represent the cluster average. Shaded contours represent the uncertainty in the motion profile. Spurious goals going off the road or in the direction of oncoming traffic are discarded.

direction. In the last scenario, the model outputs fall in both boundaries of the same lane. In this case we only output one intention. Note that this is fundamentally different from a heatmap representation, where the output is a dense grid of probabilities that assigns the likelihood of the agent being at each position on the map at a given time. While heatmaps can provide valuable information about the distribution of possible trajectories, they do not provide a clear indication of the underlying high-level intentions of the agent. This is a fundamental limitation when it is essential to disentangle different behavioural modes and predict the agent's intended goal or destination accurately. By contrast, our proposed clustering approach allows us to identify the most probable high-level intentions of the agent, providing a more interpretable and actionable output.

5.3 | Quantitative evaluation

This method not only disambiguates the multi-modal output, showing the possible intentions predicted by the model, but also improves the compliance of the output predictions. The clustering discards non-plausible goals that violate traffic rules, going off-road or in the direction of oncoming traffic.

To assess the veracity of this statement, we provide measures of DAC and OTD for both DenseTNT and our method, which we call DenseTNT-Intent.

We compute two metrics of diversity: p-RF and variance of the final heading for the different output modes, σ_{yaw}^2 for the whole output—12 modes—and the 3 most probable outputs. We propose to modify RF and compute instead the ratio of the probabilistic versions of avgFDE and minFDE, p-RF. p-avgFDE is computed following the same heuristic as p-minFDE. For each sample with K trajectories:

$$\text{p-avgFDE} = \frac{1}{K} \sum_k \text{FDE}_k + \min(-\log p_k, -\log 0.05) \quad (7)$$

Finally, we also provide probabilistic precision metrics to compare both models in terms of accuracy. All metrics must be considered simultaneously in order to get a holistic interpretation. The results are illustrated in Table 4.

DenseTNT diversity metrics are better when evaluated with the whole output, that is, 12 predictions. This is probably due to the fact that clustering removes bad predictions, which makes the variance of the yaw lower. Following the same intuition, avgFDE is higher for DenseTNT whereas minFDE is lower. Therefore, the ratio RF is also higher for DenseTNT.

TABLE 4 Quantitative results for DenseTNT (12 modes) and DenseTNT-intent in terms of diversity, admissibility, and accuracy. Improvements are indicated by arrows. The first two rows show results for the whole validation set. The third and fourth row consider top 3 predictions. The last two rows describe the results for the subset of scenes with more than one plausible mode.

Evaluation method	Model	Accuracy				Diversity		Admissibility	
		p-minADE (↓)	p-minFDE (↓)	p-avgFDE (↓)	p-MR (↓)	p-RF (↑)	σ_{yaw}^2 (↑)	DAC (↑)	OTD (%) (↓)
Whole	DenseTNT	2.71	3.03	8.03	0.79	2.65	0.096	0.97	6.99
$k = 12$	DenseTNT-Intent	1.67	3.07	4.45	0.60	1.45	0.034	0.99	0
Whole	DenseTNT	2.25	3.25	4.67	0.73	1.43	0.020	0.021	6.98
$k = 3$	DenseTNT-Intent	1.67	3.09	4.45	0.60	1.45	0.037	0.036	0
Subset	DenseTNT	2.32	3.39	4.93	0.73	1.44	0.020	0.98	7.08
$k = 12$	DenseTNT-Intent	1.89	3.28	5.43	0.68	1.66	0.058	0.98	0

Note: Values in bold represent the best performance results within the comparative analysis.

However, RF only measures diversity without looking at the quality of this output. Moreover, 90.7% of the target agent's manoeuvres go in a straight line. In most of these cases, all the uncertainty is in the motion profile, not in the intention, thus having only one output cluster. Figure 2d shows a clear example where the variance will be much higher for DenseTNT before the clustering.

We decide to evaluate this hypothesis with two more evaluation scenarios. First, we evaluate the model in the case where we only keep the three most probable predictions ($k = 3$). In this scenario, the clustered output shows higher diversity both in terms of σ_{yaw}^2 and p-RF. Secondly, we evaluate it in a subset of the validation set, considering only those scenes whose underlying distribution has more than one plausible mode. As expected, we encounter best results in terms of diversity in this evaluation scenario. For the analysis of accuracy and admissibility results, we will focus on the whole dataset with $K = 12$ since there are no major changes in the other dimensions.

The admissibility metrics show that clustering the output trajectories into high-level behaviours provides a more compliant prediction, with no clusters going in the direction of oncoming traffic and a 2pp higher DAC. We find that almost 7% of the predictions go in the opposite direction.

In order to get a holistic evaluation, we also explore how the clustering affects the results in terms of accuracy. When evaluating the probabilistic versions of distance-based metrics, DenseTNT-Intent seems to be superior in terms of p-minADE, p-avgFDE and p-MR. This is due to the fact that it provides a lower amount of modes with a higher probability for the ground-truth intention. However, since DenseTNT outputs more trajectories with higher variability in the motion profile, it is more likely that one of these trajectories will be closer to the ground-truth prior to the clustering, thus having a lower p-minFDE for DenseTNT.

These results suggest that DenseTNT-Intent achieves a more scene-compliant output, in the form of intention prediction, covering the modes of the output distribution while improving the overall quality of the predictions. This output can be used to improve the safety and reliability of autonomous vehicles by enabling better decision-making and more accurate prediction of the agents' future behaviour.

6 | INTERPRETING MULTI-MODAL MOTION PREDICTION OUTPUTS

In this section, we test our hypothesis that appropriate visualisation of multi-modal trajectory representation increases the interpretability of the prediction system. We verified our null hypothesis by creating a survey in which we explore how the visualisations of multi-modal predictions, with their associated probabilities and different number of modes and prediction horizons, impacts the interpretability of the prediction system output, and how this, consequently, might impact transparency and user trust in the overall system.

6.1 | Methodology

We recruited 39 technical experts to participate in a twenty-minute survey designed within-subjects. Twenty-three percent of the subjects are women, which is similar to the inherent bias in the technical field. In this survey, technical interpretability was assessed using different visualisation experiments of the information provided by the prediction systems.

The survey is divided into two sections. The visualisations from the first section are generated with the predictions of PGP on the nuScenes Trajectory Prediction dataset. The second section is generated with predictions of DenseTNT on Argoverse Motion Forecasting dataset. We choose the two most widely used datasets in the field of multi-modal motion prediction and the two models that ranked at the top of their benchmarks. This was done to increase the diversity of the study. In addition, it introduces differentiation in terms of the types of multi-modal traffic prediction frameworks, as DenseTNT utilises target set prediction and PGP employs lane graph sampling.

Each experiment in the survey is explored with three types of questions. First, an A/B test, in which the subject must choose between two presented cases. Second, a Likert scale, in which a five-point scale is used to allow the individual to express how much they agree or disagree with a particular statement. Finally, an open-ended question where they can explain the reason for their answers. This final question helps us to see things from the respondent's perspective, as we get

feedback in their own words. Please, refer to the supplementary materials for more details.³ In the following schema, we describe the structure of the survey and the variables under study:

- 1) Section 1. PGP visualisations in nuScenes dataset.
 - (a) Introduction and baseline visualisation with no predictions.
 - (b) Case 1. Single mode versus multi-mode.
 - (c) Case 2. Number of modes: 10 modes versus 3 modes.
 - (d) Case 3. Prediction time horizon: 6 versus 2s.
 - (e) Case 4. Importance of cameras view as an additional visualisation.
- 2) Section 2. DenseTNT and DenseTNT-intent visualisations in the Argoverse dataset.
 - (a) Introduction and baseline visualisation with no predictions.
 - (b) Case 5. Single mode versus multi-mode.
 - (c) Case 6. Trajectory prediction versus intention prediction.

These cases are designed taking into consideration the factors explored in this work, as well as those aspects that are not consistent among different evaluation benchmarks.

6.2 | Results

Table 5 details the descriptive and inferential statistics. The 5-point Likert scale questions results are analysed as interval data, being described in terms of means, variances, and mode. Parametric analysis of ordinary averages of Likert scale data is justifiable by the Central Limit Theorem. We choose Welch's Test to determine the validity of our hypothesis and determine the statistical significance of the difference in means for both distributions in each of the cases under study. Test statistic, p -value, and degrees of freedom are described in Table 5. The pie charts for the A/B type questions are shown in Figure 3. In the following, we detail our findings in each of the surveyed scenarios.

1. Single mode versus multi-mode predictions. In this first scenario, evaluated with PGP model on nuScenes dataset, we test our hypothesis that multi-modality increases the interpretability of the prediction system output over single-mode predictors. 20% of participants preferred single mode over multi-mode prediction representation in terms of interpretability of the predictive system. These subjects report finding the multi-modal visualisations more confusing, yet more reliable for the system.

We found a statistically significant difference of 0.8 points in the interpretability scoring for multi-modal predictions versus single-mode predictions, with a p -value of $9 \cdot 10^{-5}$. This is consistent with our hypothesis that multi-modality increases interpretability of the output of prediction systems.

2. Three modes versus ten modes. In this scenario, we study the optimal number of modes for better understanding of the predictive system. 41% of participants chose 10 modes over 3 modes in terms of interpretability. The most frequent argument is that the interpretation of 10 predictions may be complex for humans, but seems more reliable for the system, increasing the confidence in it. When asking to rate both cases in terms of interpretability in a 5-point Likert scale, we found no significant difference in the means of both distributions.

3. Prediction time horizon: 6s versus 2s. In this scenario, we study the importance of long-term prediction over short-term prediction for the interpretability of the output. 84.6% of the participants preferred 6 seconds of prediction over 2 seconds in terms of interpretability of the output predictions. Most argue that 6s provides more information in order to understand the scenario.

In the Likert scale rating, we found a statistically significant difference of 1.5 points in the means of both distributions, with a p -value of 10^{-10} . This is consistent with our hypothesis that longer predictions horizons increase the interpretability of the predictions output and the reliability in the system.

4. Cameras view as additional visualisation. In this scenario, we explored the influence of adding the visualisation of six cameras placed in the AV on the interpretability of the traffic scene. 70% of participants believe the additional visualisation helps them understand the scene. 87.2% believe that it is easier to understand the behaviour of the focal agent and its predicted trajectory.

5. Single mode versus multi-mode predictions with DenseTNT on the Argoverse dataset. We repeat the first scenario on a different setup to evaluate the difference in scene representation of the Argoverse and nuScenes datasets, as well as the output representation of DenseTNT versus PGP model. We found different results from those of the first scenario. In this case, 33.3% of participants chose single mode predictions over multi-mode predictions in terms of interpretability. The biggest difference, however, is in the Likert scale rating. In this case, both means are almost identical, with no statistical difference between distributions. This notable difference is mainly due to the output representation in the case of DenseTNT, which outputs 12 trajectories with no probability associated. People believe that this makes it difficult to interpret the output. This result is yet another justification for the need of a posterior clustering into intentions with associated probabilities, which leads us to the last scenario.

6. Trajectory prediction versus intention prediction. We test the hypothesis presented in Section 5. We posit that intention prediction can contribute towards the interpretability of the system since it disambiguates the multi-modal prediction output and produces a clustered output that is more in line with the way human drivers think. 87.2% of participants believe the clustered output with its associated probability is more understandable when visualising the future predicted behaviour of the focal agent.

Most argument argue that they find it more useful to know the high-level behaviour than the concrete future trajectory, as

³Link to the survey.

	$k = 1$ versus $k > 1$ (nS)		$k = 3$ versus $k = 10$		6 versus 2s		$k = 1$ versus $k > 1$ (AV)		Trajectories versus intentions	
	$k = 1$	$n > 1$	$k = 3$	$k = 10$	6s	2s	$k = 1$	$n > 1$	Trajectories	Intentions
Mean	3.410	4.179	3.821	3.641	4.333	2.872	3.282	3.589	3.205	4.102
Var	0.879	0.458	0.835	1.025	0.3333	0.6932	0.997	1.143	0.693	0.568
Mode	3	4	4	4	4	3	3	4	3	4
t-Stat	-4.139		0.822		9.007		-1.313		-4.989	
p-value	$9 \cdot 10^{-5}$		0.41		10^{-10}		0.193		$1.9 \cdot 10^{-6}$	
df	69		75		67		75		75	

TABLE 5 Evaluation of survey results: Descriptive and inferential statistics for each of the cases under study. k is the number of modes, being $n > 1$ the multi-modal scenario. We provide mean, variance, and mode as descriptive statistics. For inferential statistics, test statistic, p value and degrees of freedom (dg) based on Welch's test are shown. nS refers to nuScenes, AV refers to the Argoverse dataset.

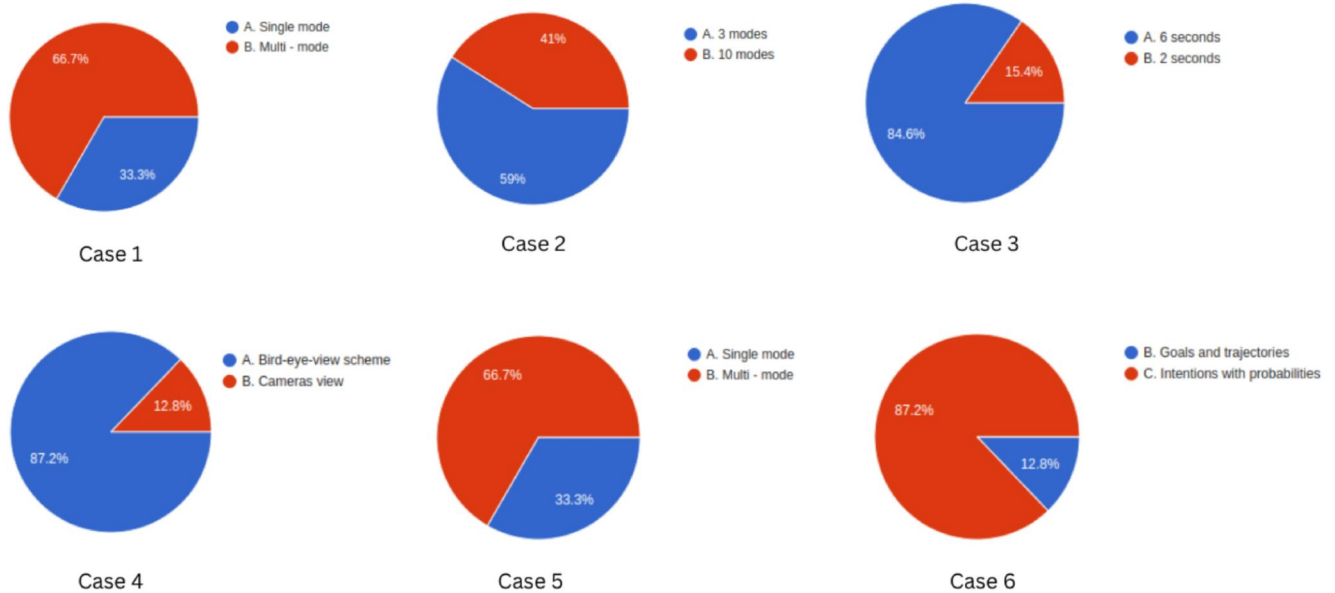


FIGURE 3 A/B questions pie charts.

it is more in line with the way humans reason. Others pose it is faster to interpret and adds more information while removing the ambiguity of multi-modal trajectory prediction. In addition, they also indicate the importance of having uncertainties associated with the predictions, to allow the subsequent algorithm to adopt a probabilistic framework. However, 5 participants believe that the second option is more confusing and difficult to understand.

When looking at the inferential statistics of the Likert scale question, there is a statistically significant difference of 0.9 points in the mean of both distributions, with a p -value of $3 \cdot 10^{-6}$. This verifies our hypothesis that intention prediction helps improve the understanding of the predictive system.

6.3 | Discussion

This study comes with several limitations that should be considered. First, Likert scales are subject to distortion due to *central tendency bias* (i.e., avoidance of using extreme categories), and *acquiescence bias* (i.e., agreeing with statements in the survey).

Second, the A/B type questions leave no room for a neutral response when no option is preferred. However, this could be specified in the Likert scale and in the open-ended question of the particular scenario. It has been taken into consideration for the analysis.

Finally, some subjects claimed to be unsure of the purpose of some questions and believe that some answers are highly dependent on the particular task. For example, for a developer it may be more interesting to have more information at the cost of a more complex representation. Some questions also depend on the traffic scenario, speed, and complexity of the road, with the optimal number of modes or prediction horizon being different for each case.

Despite these limitations, we can derive some conclusions from this study.

First, it verifies our hypothesis that multi-modality is not only important for better decision making and planning a safer route, but it also improves the understanding of prediction systems when properly visualised. Another important insight is that the form of visualisation is crucial and it is important to show the different probabilities associated with each trajectory in a clear way.

Second, longer prediction time horizons provide a better understanding of the whole traffic scene and make the overall system more interpretable.

Finally, the need for intention prediction is supported by the survey findings. Clustered outputs provide a more human-like representation of the predictions into high-level behaviours.

7 | CONCLUSIONS

In this work, we move towards the design of reliable motion prediction models based on evaluation, robustness, and interpretability of the outputs.

Our findings can be summarised as follows:

- We highlighted the main gaps and differences in current evaluation methodologies, especially in terms of lack of diversity assessment and admissibility with the traffic scene. We identify the main aspects that are critical for the evaluation of multi-modal motion prediction and propose a more comprehensive and holistic evaluation framework.
- In the robustness analysis, we showed how failure to perceive the road topology has a greater impact on system performance compared to failure to perceive other agents on the road, due to the significant inductive biases introduced by the lanes. This also showcases the need for more comprehensive datasets covering complex scenes with high interaction and a wide range of edge cases.
- We provided DenseTNT-intent outputs with high-level intentions that prove to be diverse, compliant, and accurate, improving the overall quality of the predictions.
- The results of the study suggest that this new representation improves the interpretability of the output prediction. Our first hypothesis that multi-modality improves interpretability over single-mode predictions is also verified. Finally, long-term predictions appear to provide a better understanding of the predicted traffic scene.

The proposed approach and findings make a significant contribution to the development of trustworthy motion prediction systems for autonomous vehicles. By comprehensively analysing current evaluation metrics, identifying gaps, and proposing a new holistic evaluation framework, this work provides a valuable foundation for future research in the field. Additionally, the formulation of a method for assessing spatial and temporal robustness, as well as the proposed intent prediction layer, demonstrate innovative solutions for addressing the complex challenges of multi-modal motion prediction. Finally, the assessment of interpretability through a survey of different visualisation techniques offers further insights for enhancing the performance and transparency of autonomous vehicle systems. Overall, this work represents a substantial step forward in the design of trustworthy artificial intelligence for safe and efficient autonomous driving.

ACKNOWLEDGEMENTS

This work was supported by HUMAINT project at the Algorithmic Transparency Unit at the Directorate-General Joint Research Centre (JRC) of the European Commission, and in part by the Spanish Ministry of Science and Innovation under Grant PID2020-114924RB-I00, and by the Community Region of Madrid under Grant S2018/EMT-4362 SEGVAUTO 4.0-CM. In addition, part of this work has been carried out during the *Eye for AI Program* thanks to the support of Zenseact and AI Sweden.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.



DISCLAIMER

The information and views expressed in this paper are purely those of the authors and do not necessarily reflect an official position of the European Commission.

ORCID

Sandra Carrasco Limeros  <https://orcid.org/0000-0002-4782-7166>

Sylwia Majchrowska  <https://orcid.org/0000-0001-7576-7167>

Joakim Johnander  <https://orcid.org/0000-0003-2553-3367>
 Christoffer Petersson  <https://orcid.org/0000-0002-9203-558X>

Miguel Ángel Sotelo  <https://orcid.org/0000-0001-8809-2103>

David Fernández Llorca  <https://orcid.org/0000-0003-2433-7110>

REFERENCES

1. Rudenko, A., et al.: Human motion trajectory prediction: a survey. *Int. J. Robot. Res.* 39(8), 895–935 (2020). <https://doi.org/10.1177/0278364920917446>
2. Biparva, M., et al.: Video action recognition for lane-change classification and prediction of surrounding vehicles. *IEEE Transactions on Intelligent Vehicles* 7(3), 569–578 (2022). <https://doi.org/10.1109/tiv.2022.3164507>
3. Fernandez-Llorca, D., Gómez, E.: Trustworthy artificial intelligence requirements in the autonomous driving domain. *Computer* 56(2), 29–39 (2023). <https://doi.org/10.1109/mc.2022.3212091>
4. European Commission.: Content Directorate-General for Communications Networks, and Technology. Ethics guidelines for trustworthy AI. Publications Office (2019)
5. European Commission.: Content Directorate-General for Communications Networks, and Technology. The Assessment List for Trustworthy Artificial Intelligence (ALTAI) for Self Assessment. Publications Office (2020)
6. EC.: Regulation of the European Parliament and the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts (2021). COM(2021) 206 final. <https://eur-lex.europa.eu/legal-content/EN/TEXT/?uri=CELEX>

7. Office of U.S. Senator Ron Wyden: Algorithmic Accountability Act of 2022. 117th Congress 2D Session (2022). <https://www.congress.gov/bill/117th-congress/house-bill/6580/text>
8. Chang, M.-F., et al.: Argoverse: 3D tracking and forecasting with rich maps. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 8748–8749 (2019)
9. Caesar, H., et al.: nuScenes: a multimodal dataset for autonomous driving. In: Computer Vision and Pattern Recognition (CVPR), pp. 11621–11631 (2020)
10. Gu, J., Sun, C., Zhao, H.: Densetnt: end-to-end trajectory prediction from dense goal sets. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 15283–15292 (2021)
11. Deo, N., Wolff, E., Beijbom, O.: Multimodal trajectory prediction conditioned on lane-graph traversals. In: 5th Annual Conference on Robot Learning (2021)
12. Tang, Y.C., Salakhutdinov, R.: Multiple futures prediction. In: 33rd Conference on Neural Information Processing Systems (NeurIPS 2019) (2019)
13. Tim Salzmänn, et al.: Trajectron++: dynamically-feasible trajectory forecasting with heterogeneous data. In: European Conference on Computer Vision (ECCV) (2020)
14. Konev, S., Brodt, K., Sanakoyeu, A.: Motioncnn: a strong baseline for motion prediction in autonomous driving. In: Workshop on Autonomous Driving, CVPR (2021)
15. Liu, Y., et al.: Multimodal motion prediction with stacked transformers. In: Computer Vision and Pattern Recognition (CVPR), pp. 7577–7586 (2021)
16. Varadarajan, B., et al.: Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction. In: 2022 International Conference on Robotics and Automation (ICRA), pp. 7814–7821 (2022)
17. Gilles, T., et al.: THOMAS: trajectory heatmap output with learned multi-agent sampling. In: International Conference on Learning Representations (2022)
18. Gao, J., et al.: VectorNet: encoding HD maps and agent dynamics from vectorized representation. In: Computer Vision and Pattern Recognition (CVPR), pp. 11525–11533 (2020)
19. Zhao, H., et al.: Tnt: target-driven trajectory prediction. In: Volume 155 of *Proceedings of Machine Learning Research*, pp. 895–904. PMLR (2021)
20. Carrasco, S., Fernández Llorca, D., Sotelo, M.A.: SCOUT: socially-consistent and understandable graph attention network for trajectory prediction of vehicles and VRUs. In: 2021 IEEE Intelligent Vehicles Symposium (IV), pp. 1501–1508 (2021)
21. Cui, H., et al.: Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In: 2019 International Conference on Robotics and Automation (ICRA), pp. 2090–2096 (2019)
22. Casas, S., et al.: SpAGNN: spatially-aware graph neural networks for relational behavior forecasting from sensor data. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), pp. 9491–9497 (2020)
23. Wang, J., Tengyu Ye, Z.G.1, Chen, J.: LTP: lane-based trajectory prediction for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 17134–17142 (2022)
24. Khandelwal, S., et al.: What-if motion prediction for autonomous driving. ArXiv, abs/2008.10587 (2020)
25. Gilles, T., et al.: GOHOME: graph-oriented heatmap output for future motion estimation. In: 2022 International Conference on Robotics and Automation (ICRA), pp. 9107–9114 (2021)
26. Liang, M., et al.: Learning lane graph representations for motion forecasting. In: European Conference on Computer Vision (ECCV) (2020)
27. Zeng, W., et al.: LaneRCNN: distributed representations for graph-centric motion forecasting. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 532–539 (2021)
28. Cho, K., et al.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. In: 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1724–1734 (2014)
29. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* 9(8), 1735–80 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
30. Vaswani, A., et al.: Attention is all you need. In: Guyon, I., et al. (eds.) *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc. (2017)
31. Luo, C., et al.: Probabilistic multi-modal trajectory prediction with lane attention for autonomous vehicles. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 2370–2376 (2020)
32. Phan-Minh, T., et al.: Covernet: multimodal behavior prediction using trajectory sets. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020). –14071.14062
33. Chai, Y., et al.: MultiPath: multiple probabilistic anchor trajectory hypotheses for behavior prediction. In: 3rd Conference on Robot Learning (CoRL) (2019)
34. Makansi, O., et al.: Overcoming limitations of mixture density networks: a sampling and fitting framework for multimodal future prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
35. Cao, D., et al.: Spectral temporal graph neural network for trajectory prediction. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 1839–1845 (2021)
36. Gupta, A., et al.: Social gan: socially acceptable trajectories with generative adversarial networks. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2255–2264 (2018)
37. Lee, N., et al.: Desire: distant future prediction in dynamic scenes with interacting agents. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2165–2174 (2017)
38. Kim, S., et al.: Improving diversity of multiple trajectory prediction based on map-adaptive lane loss. ArXiv, abs/2206.08641 (2022)
39. Scott, E., 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 (ICCV), pp. 9710–9719 (2021)
40. Casas, S., et al.: Implicit latent variable model for scene-consistent motion forecasting. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
41. Brier, G.W.: Verification of forecasts expressed in terms of probability. *Mon. Weather Rev.* 78(1), 1–3 (1950). [https://doi.org/10.1175/1520-0493\(1950\)078<0001:vofeit>2.0.co;2](https://doi.org/10.1175/1520-0493(1950)078<0001:vofeit>2.0.co;2)
42. Ivanovic, B., Pavone, M.: The trajectron: probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In: International Conference on Computer Vision (ICCV), pp. 2375–2384 (2019)
43. Park, S.H., et al.: Diverse and admissible trajectory forecasting through multimodal context understanding. In: 16th European Conference on Computer Vision (ECCV) (2020)
44. Hu, Y., Zhan, W., Tomizuka, M.: Probabilistic prediction of vehicle semantic intention and motion. In: 2018 IEEE Intelligent Vehicles Symposium (IV), pp. 307–313 (2018)
45. Casas, S., Luo, W., Urtasun, R.: IntentNet: learning to predict intention from raw sensor data. In: In 2nd Conference on Robot Learning (CoRL) (2018)

How to cite this article: Carrasco Limeros, S., et al.: Towards trustworthy multi-modal motion prediction: holistic evaluation and interpretability of outputs. *CAAI Trans. Intell. Technol.* 1–16 (2023). <https://doi.org/10.1049/cit2.12244>