

# Multi-modal Motion Prediction using Temporal Ensembling with Learning-based Aggregation

Kai-Yin Hong, Chieh-Chih Wang and Wen-Chieh Lin

**Abstract**—Recent years have seen a shift towards learning-based methods for trajectory prediction, with challenges remaining in addressing uncertainty and capturing multi-modal distributions. This paper introduces *Temporal Ensembling with Learning-based Aggregation*, a meta-algorithm designed to mitigate the issue of missing behaviors in trajectory prediction, which leads to inconsistent predictions across consecutive frames. Unlike conventional model ensembling, temporal ensembling leverages predictions from nearby frames to enhance spatial coverage and prediction diversity. By confirming predictions from multiple frames, temporal ensembling compensates for occasional errors in individual frame predictions. Furthermore, trajectory-level aggregation, often utilized in model ensembling, is insufficient for temporal ensembling due to a lack of consideration of traffic context and its tendency to assign candidate trajectories with incorrect driving behaviors to final predictions. We further emphasize the necessity of learning-based aggregation by utilizing mode queries within a DETR-like architecture for our temporal ensembling, leveraging the characteristics of predictions from nearby frames. Our method, validated on the Argoverse 2 dataset, shows notable improvements: a 4% reduction in minADE, a 5% decrease in minFDE, and a 1.16% reduction in the miss rate compared to the strongest baseline, QCNet, highlighting its efficacy and potential in autonomous driving.

**Index Terms**—Autonomous driving, multi-modal motion prediction, DETR, ensembling.

## I. INTRODUCTION

Autonomous driving technology, since its inception in the 1980s with Pomerleau’s groundbreaking work [1], has profoundly impacted daily lives, particularly in ensuring safe and comfortable path planning and collision avoidance through accurate anticipation of surrounding traffic. Motion prediction, therefore, plays a crucial role in various autonomous driving applications, including risk estimation [2], decision making [3], and traffic simulation [4]. In recent years, motion prediction has increasingly relied on learning-based methods [5]–[10]. Despite significant progress in this field, substantial challenges persist, particularly in capturing multi-modality and dealing with the considerable uncertainty in output predictions, especially as the prediction horizon extends. For instance, when a vehicle approaches an intersection, its actions may vary depending on the driver’s long-term




Fig. 1: Illustration of Missing Behaviors Issue - Missing behaviors refer to cases where predictions are occasionally wrong, resulting in inconsistent predictions across consecutive frames. The left panel depicts multi-modal motion prediction. The red car represents the target agent, with its red trajectory showing the ground truth. The gray trajectories illustrate various possible future paths. The middle and right panels demonstrate missing behaviors in the predicted trajectories across consecutive time steps.

goals, leading to multiple possible trajectories. Overcoming this challenge requires motion prediction models to learn and capture the underlying multi-modal distribution instead of predicting only the most common mode. However, this task is complex, as each training sample typically represents only one possibility.

When applying the state-of-the-art approach, QCNet [10] to real-world scenarios, we observed instances where predictions may be inaccurate and fail to capture the exact behavior. For example, when examining the predicted trajectories at time step  $t$  in Fig. 1, all predictions suggest a left turn, but the actual movement turns out to be a right turn. Furthermore, upon closer observation, we observed that the predictions at time step  $t - 2$  were accurate, despite the error at time step  $t$ . We aim to leverage this characteristic to address missing behaviors.

We first propose **Temporal Ensembling** to address missing behaviors. By integrating predictions from nearby frames, we aim to compensate for occasional errors in individual frame predictions, enhance spatial coverage, and improve accuracy. Similar to the model ensembling employed by [8], [10]–[12] where multiple models are trained and their predictions combined to enhance performance, temporal ensembling accomplishes motion predictions with a single model but with predictions from multiple frames.

Temporal ensembling shares similarities with model ensembling by increasing the pool of prediction candidates before aggregating them into the final predictions. However, trajectory-level aggregation methods such as Top-K, Non-Maximum Suppression (NMS), and K-means - commonly used in model ensembling [8], [10], [11] - are found to be inadequate for temporal ensembling. This is because trajectory-level aggregation does not consider traffic con-

Kai-Yin Hong is with the Institute of Electrical and Computer Engineering, National Yang Ming Chiao Tung University, Hsinchu, Taiwan. kaiyin0208.ee11@nycu.edu.tw

Chieh-Chih Wang is with the College of Electrical and Computer Engineering, National Yang Ming Chiao Tung University, and with the Mechanical and Mechatronics Systems Research Laboratories, Industrial Technology Research Institute, Hsinchu, Taiwan. bobwang@ieee.org

Wen-Chieh Lin is with the Institute of Multimedia Engineering, National Yang Ming Chiao Tung University, Hsinchu, Taiwan wclin@cs.nctu.edu.tw

text and tends to incorrectly assign candidate trajectories, which exhibit inappropriate driving behaviors, to the final predictions. This limitation motivates the need for a more dynamic and context-aware approach. We propose **Learning-based Aggregation** to enable temporal ensembling by utilizing mode queries within a DETR-like architecture [13] to determine the final predictions.

The main contribution of this paper is introducing a meta-algorithm named Temporal Ensembling with Learning-based Aggregation to address missing behaviors. We validated our approach on the Argoverse 2 dataset [14], achieving notable improvements in the three crucial prediction metrics: a 4% enhancement in minADE, a 5% improvement in minFDE, and a 1.16% reduction in miss rate compared to the strongest baseline, QCNet [10].

## II. RELATED WORKS

Predicting vehicle behavior is crucial for autonomous driving. Initially, research focused on physics-based models for short-term predictions, as highlighted in studies such as [2] and [15]. Recently, the trend has shifted to learning-based methods, which are preferred for their accuracy and ability to account for road interactions [16]. We’ll review the latest work in multi-modal motion prediction and trajectories ensembling techniques. For a complete review on motion prediction, please see [16].

### A. Multi-Modal Motion Prediction

Due to the uncertainty in agent intent, motion prediction outputs naturally exhibit multiple modes, making it challenging to determine whether a vehicle will proceed straight or turn right as it approaches an intersection. Therefore, having a model capable of capturing multiple potential trajectories within a limited set is crucial. Several studies [17]–[19] aim to predict multiple future trajectories and associated probabilities through direct regression. To enhance the coverage of potential outcomes, anchor-based approaches have gained popularity in trajectory prediction. These approaches first classify discrete intent and then regress continuous trajectories conditioned on the identified intent. Intent classification can be categorized into two types: goal-based (goal positions [20], [21] or target lanes [18]) and driving maneuver-based [22]. The recent state-of-the-art method [10] utilizes the DETR-like [13] architecture to address the multi-modal problem, combining anchor-free and anchor-based techniques to achieve notable performance. Despite advancements in multi-modal problems, a gap remains in achieving both accurate and comprehensively covered predictions. Our analysis of the existing model [10] identifies inherent limitations, termed missing behaviors, as illustrated in Fig. 1. This insight has led us to develop strategies to mitigate the multi-modal problem, focusing on tackling the issue of missing behaviors.

### B. Trajectories Ensembling Techniques

Model ensembling [23] is a widely adopted technique to enhance motion prediction performance [8], [10]–[12]. Typical strategies involve training multiple replicas of the model




Fig. 2: Comparison of Ensembling Methods - The figure illustrates two ensembling approaches. **Model Ensembling (Top)**: Multiple models independently predict  $N$  trajectories at Frame  $t$ . With  $M$  models, this results in  $M * N$  trajectories that are combined into the final  $N$  trajectories at the trajectory level. **Temporal Ensembling (Bottom)**: A single model generates  $M * N$  predictions across  $M$  nearby frames. Our proposed learning-based aggregation then combines them into the final  $N$  trajectories.

with varying initial parameters, learning rates, epochs, or data portions. In these model ensembling-based approaches, while a single model generates  $N$  trajectories, employing  $M$  models can yield  $M * N$  trajectories. However, post-processing is necessary to integrate outputs from multiple models for a fair comparison with the benchmark. Seminal works in this domain include MultiPath++ [8], which combines Top- $K$  trajectories based on confidence scores to integrate multiple model results; ProphNet [11] utilizes Non-Maximum Suppression (NMS) and QCNet [10] leverages K-means for integrating multiple model results, respectively. Top- $K$  selects the top  $K$  trajectories based on confidence scores. NMS combines geometric constraints with the Top- $K$  to balance high scores and diversity by selecting trajectories with high scores while excluding similar ones. K-means clusters trajectories based on the endpoints and performs a weighted average of trajectories within each cluster according to confidence scores. Our proposed temporal ensembling, while similar to model ensembling in its method of expanding the pool of prediction candidates, distinguishes itself through learning-based aggregation. Our proposed aggregation is achieved through mode queries within a DETR-like architecture [13] and considers traffic context, rather than directly aggregating predictions at the trajectory level.

## III. PROPOSED METHOD

This section starts with the baseline utilized and introduces temporal ensembling to address the issue of missing behaviors. It then explains how naively applied trajectory-level aggregation to temporal ensembling is insufficient, leading to the proposal of a learning-based aggregation method.

### A. Baseline Model

1) *Input Output Formulation*: QCNet [10] was selected as our baseline model due to its state-of-the-art performance. QCNet is the winner of the Argoverse Motion Forecasting 2023 competition. It employs an encoder-decoder architecture with vectorized representation [5] as input. The agent state comprises the spatial position  $p_i^t = (p_{i,x}^t, p_{i,y}^t)$ , angular position  $\theta_i^t$ , temporal time  $t$ , and motion vector  $p_i^t - p_i^{t-1}$  for the  $i$ -th agent at time step  $t$ , while the HD map contains spatially sampled points and semantic attributes. Input coordinates are set in the world coordinate system while the output trajectories are in the agent-centric coordinate system.

This design eliminates redundant input encoding associated with agent-to-agent centric mapping [9], [11], leading to significant savings in computational resources.

2) *Encoder*: To achieve the world-to-agent coordinate transformation, the encoder utilizes polar representation within each scene element and relative spatial-temporal positional embedding [10] between every pair of scene elements. Subsequently, the Fourier transform [24] is applied to scene elements, followed by Multi-Layer Perceptron (MLP) projection to higher-dimensional space. Finally, a factorized attention approach [7], [9], [10] is employed to fuse different entities, resulting in the ultimate scene embedding.

3) *Decoder - DETR-like Architecture*: Recently, trajectory prediction methods [10], [11], [25], [26] inspired by DETR [13] have gained popularity. The DETR-based approach is particularly effective in addressing one-to-many problems, where a single scene embedding corresponds to predicting multiple trajectories. This is achieved through the design of mode queries and the Transformer [27] structure, involving multiple learnable queries that cross-attend the scene embedding and decode trajectories. For more architectural details, refer to QCNet [10].

4) *Training Objectives*: Following the approaches [9], [10], we parameterize the future trajectory of the  $i$ -th agent as a mixture of Laplace distributions:

$$f(\{p_i^t\}_{t=1}^{T'}) = \sum_{n=1}^N \pi_{i,n} \prod_{t=1}^{T'} \text{Laplace}(p_i^t | \mu_{i,n}^t, b_{i,n}^t), \quad (1)$$

where  $\{\pi_{i,n}\}_{n=1}^N$  represents the mixing coefficients, and the Laplace density of the  $n$ -th mixture component at time step  $t$  is defined by the parameters  $\mu_{i,n}^t$  and  $b_{i,n}^t$ . A classification loss  $\mathcal{L}_{cls}$  is utilized to optimize the mixing coefficients. This loss minimizes the negative log-likelihood of Eq. 1. On the other hand, a winner-take-all strategy [28] is applied to  $\mathcal{L}_{traj}$ , conducting backpropagation solely on the best-predicted trajectory. The total loss function for end-to-end training is given by:

$$\mathcal{L}_{total} = \mathcal{L}_{traj} + \lambda \mathcal{L}_{cls}, \quad (2)$$

where  $\lambda$  balances regression and classification. Optimization is carried out using the AdamW optimizer [29] over 64 epochs, with a batch size of 32, a dropout rate of 0.1, and a weight decay coefficient of  $1 \times 10^{-4}$ . The initial learning rate is set to  $5 \times 10^{-4}$  and decayed using the cosine annealing scheduler.

## B. Temporal Ensembling

### 1) Naive Approach with Trajectory-level Aggregation:

In our initial approach to integrating trajectories predicted across multiple time steps, as shown in Fig. 2, we utilize the property of large overlaps, depicted in Fig. 5. It is found that when evaluating the required horizon at time step  $t$ , predictions can be made as early as at time step  $t - M + 1$ . This recognition of overlap in prediction horizons allows for the confirmation of predictions across frames, compensating for errors in single-frame predictions. Initially, trajectories

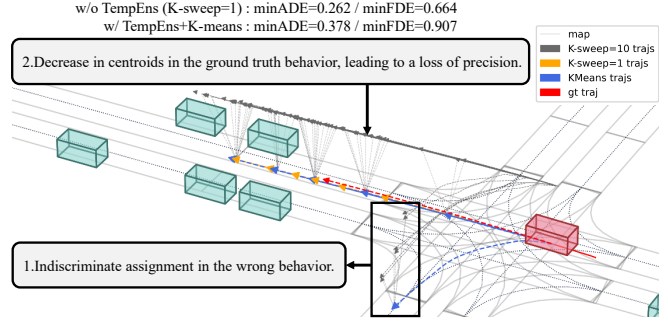


Fig. 3: This highlights the precision-diversity trade-off in trajectory-level aggregation using K-means. Gray trajectories represent all trajectories within the sliding window. Orange trajectories depict single-frame approach predictions, while blue trajectories demonstrate the integration of multiple-frame predictions aggregated from the gray ones.

for each nearby time step are predicted. Subsequently, overlapping segments with the target horizon at time step  $t$  are sliced, ensuring within the same time range. At this point, we achieve  $M \times N$  trajectories, as shown in the middle panel of Fig. 2, with  $N$  being the number of trajectories predicted per time step and  $M$  the number of frames to integrate. A final aggregation stage then merges these into the ultimate  $M$  trajectories, illustrated on the rightmost side of Fig. 2. For thorough evaluation, we set  $M = 10$  and  $N = 6$ .

TABLE I: Performance comparison of temporal ensembling with trajectory-level aggregation techniques on the Argoverse 2 validation set. The baseline observes 50 frames, predicts 60 future frames, and evaluates at intervals [10,60]. For a fair comparison, single-frame predictions at  $t = 10$  are made to match temporal ensembling by including the most recent agent states.

Method	minADE	minFDE	MissRate
QCNet [10]	<b>0.50</b>	<b>0.99</b>	10.73%
QCNet [10] + TempEns w/ Top-K	0.72	1.65	24.80%
QCNet [10] + TempEns w/ NMS	0.64	1.19	12.27%
QCNet [10] + TempEns w/ K-means	0.51	1.01	<b>10.72%</b>

In our initial experiment, which involved integrating trajectory-level aggregation into temporal ensembling, as illustrated in Table I, we discovered that trajectory-level aggregation was less effective than anticipated. Specifically, predictions based on a single frame outperformed those derived from multiple frames. Our analysis revealed a pattern across multiple frames predictions where high-scoring trajectories often exhibit similarity. This phenomenon adversely affects when applying selection manners (Top-K and NMS), leading to a decrease in diversity and consequently degrading performance. Further analysis presented in Fig. 3 reveals two key insights regarding the clustering manner (K-means). First, integrating trajectories through distance-based clustering enhanced diversity by preserving geometrical variations, including both straight and left-turn trajectories. Second, however, this approach sometimes reduced accuracy, particularly when the ground truth was a straight trajectory. K-means' indiscriminate assignment of candidate trajectories to different behaviors, without re-considering traffic context, often resulted in inaccurate aggregation.

2) *Learning-based Aggregation*: To address the issues associated with trajectory-level aggregation, we propose in-

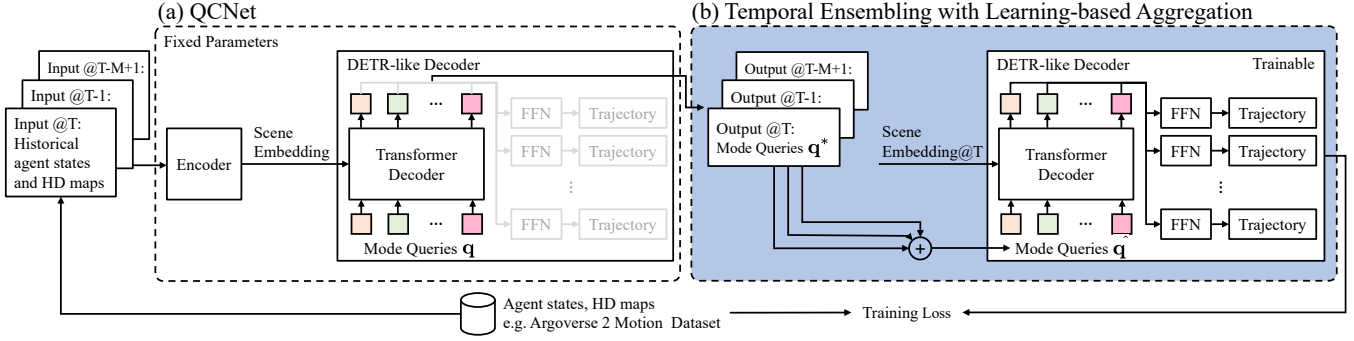


Fig. 4: **Overall Pipeline of Temporal Ensembling with Learning-based Aggregation** - The architecture consists of two main blocks. **Block (a)** represents the baseline model, QCNet, from which we leverage the predicted mode queries (not the final trajectories). **Block (b)** depicts our proposed method. It takes predictions of mode queries from nearby frames as input. Element-wise addition is used to aggregate historical mode queries. A transformer decoder then fuses the aggregated mode queries with scene embedding at time step  $T$ . Finally, each feed-forward network (FFN) predicts the final trajectory.

tegrating the aggregation process into the model pipeline, with an emphasis on considering the traffic context. We propose leveraging mode query designs in a DETR-like decoder [10], marking a shift toward using vectors formed by the transformer decoder with scene embeddings, as opposed to directly using trajectories. This vector encapsulates potential driving intentions, which are then transformed into trajectories. By focusing on this high-dimensional space, we identify an ideal opportunity for learning-based aggregation. This technique allows us to combine driving intentions across time, effectively addressing behaviors that are missed in predictions based on a single frame. As depicted in Fig. 4, the mode queries we utilized are expressed as:

$$\mathbf{q}_{nt}^* = \text{TransformerDecoder}(\mathbf{q}_{nt}, \mathbf{s}_t), \quad (3)$$

where  $\mathbf{q}_{nt}$  denotes the  $n$ -th initial mode query at time step  $t$ , with  $n \in 1, 2, \dots, N$  and  $N$  representing the total number of mode queries.  $\mathbf{s}_t$  represents the scene embedding generated by the encoder at time step  $t$ . Each  $\mathbf{q}_{nt}^*$ , the output of the transformer decoder, is tasked with predicting final trajectories and representing driving intentions within the spatial-temporal scene context. We collect historical mode queries from nearby frame predictions, each generating  $N$  mode queries, resulting in  $M \times N$  mode queries, where  $M$  is the number of nearby frames. Through element-wise addition, denoted as

$$\hat{\mathbf{q}}_{nt} = \mathbf{q}_{nt}^* + \mathbf{q}_{nt-1}^* + \dots + \mathbf{q}_{nt-M+1}^*, \quad (4)$$

we combine the corresponding positions of the  $N$  mode queries  $\mathbf{q}_{nt}^*$  to obtain the aggregated mode query  $\hat{\mathbf{q}}_{nt}$ . The slight time offset between the  $N$  mode queries  $\mathbf{q}_{nt}^*$  for nearby frames, there is a need to re-utilize the scene embedding  $\mathbf{s}_t$  to refine  $\hat{\mathbf{q}}_{nt}$ . We employ a baseline transformer decoder [10] to fuse  $\hat{\mathbf{q}}_{nt}$  with the scene embedding  $\mathbf{s}_t$ . This creates a context-aware mode query  $\tilde{\mathbf{q}}_{nt}$ . Feed Forward Networks (FFNs) then convert the final state of  $N$  mode queries  $\tilde{\mathbf{q}}_{nt}$  into  $N$  trajectories as follows:

$$\tilde{\mathbf{q}}_{nt} = \text{TransformerDecoder}(\hat{\mathbf{q}}_{nt}, \mathbf{s}_t), \quad (5)$$

$$\text{trajectory}_{nt} = \text{FFN}(\tilde{\mathbf{q}}_{nt}). \quad (6)$$

TABLE II: This table presents the results of two aggregation operations within our learning-based aggregation method on the Argoverse 2 validation set. The baseline observes 50 past frames to predict 60 future steps (evaluated at intervals [10, 60]).

Method	minADE	minFDE	MissRate
LearnAgg w/ CrossAttn	0.48	0.95	9.72%
LearnAgg w/ ElementWiseAdd	<b>0.48</b>	<b>0.94</b>	<b>9.57%</b>

### 3) Two Operations within Learning-based Aggregation:

In learning-based aggregation, the most straightforward operation involves using a cross-attention mechanism [27], where queries are the mode queries from the current frame, and keys and values are the anticipated integrated mode queries from nearby frames. Cross-attention aims to harness the capability of attention mechanisms to identify crucial behavioral features that might be missed in the current frame. Its effectiveness is shown in Table II, with improvements in all three main metrics. However, by leveraging the DETR-like structure [13] and the design of mode queries, we further explore a simpler operation: element-wise addition of mode queries. This approach is motivated by the desire to capture missing behaviors from nearby frame predictions. At a high level, it aims to encompass driving behaviors by incorporating the right behaviors into the current prediction. By overlapping the high-dimensional vectors, the full spectrum of driving intentions from nearby frames can be covered, which yields even better results. This improvement is attributed to the additive nature of mode queries.

4) *Training Pipeline*: Our proposed training pipeline for temporal ensembling with learning-based aggregation, as depicted in Fig. 4, comprises two blocks. In Fig. 4-(a), a fully trained base model is prepared, and its parameters are then frozen. Subsequently, predictions are made on continuous data. The mode queries across multiple frames produced by the base model are collected and used as the proposed inputs. These are then fed into the newly added decoder [10], as shown in Fig. 4-(b), resulting in the final set of  $N$  trajectories. The training loss for our proposed method is aligned with the loss of the base model. The parameters of the newly added decoder are fine-tuned, starting with an initial learning rate of  $2.5e-4$ , which is half of the base model. The AdamW




Fig. 5: Streaming-style formulation - Predictions exhibit a high degree of overlap in continuous datasets in the streaming-style paradigm. The overlapping time ranges, shown by the two dashed lines, offer an opportunity to exploit this property.

optimizer [29], along with a cosine annealing scheduler, is employed for a total duration of 8 training epochs.

#### IV. EXPERIMENTS

In this section, we commence by presenting comprehensive information about the dataset and metrics utilized. Subsequent comparisons with the state-of-the-art methods and qualitative results are provided. Additionally, the combined effect of model ensembling and the proposed temporal ensembling is explored to demonstrate their extensibility and effectiveness.

##### A. Experimental Setup

1) *Dataset and Streaming-Style Formulation*: Our method is evaluated using the Argoverse 2 Dataset [14] which is widely recognized for its diverse and challenging scenarios. This dataset contains 250K non-overlapping scenarios. Trajectories are sampled at 10Hz, with an observation window of  $(-50, 0]$  frames and a prediction horizon of  $(0, 60]$  frames. Each scenario is associated with a local map region, including geometric properties such as lane centerlines, boundaries, crosswalks, and road markers. We redefine the selected dataset [14] to exhibit streaming-style characteristics by transforming each data segment. Rather than predicting each snapshot once within a fixed-length segment, our approach employs a sliding-window temporal arrangement, as depicted in Fig. 5, to introduce streaming characteristics. This arrangement can be achieved through two approaches within the selected dataset [14]: One approach maintains the original configuration of the model while shortening the evaluation length. The other conforms to the benchmark evaluation setting by adjusting the model to observe shorter periods and predict beyond the evaluation length. Experiments are conducted in both settings to demonstrate the effectiveness of the proposed method.

2) *Metrics*: To evaluate prediction performance, standard metrics are employed: minimum final displacement error (minFDE), minimum average displacement error (minADE), and miss rate, aiming for lower values across up to  $N=6$  predicted trajectories per agent. **minFDE** measures the  $\ell^2$  distance between the endpoints of ground truth trajectory and predicted trajectory  $\hat{N}$ , where  $\hat{N}$  is the trajectory index with the minimum endpoint error among  $N$  predicted trajectories. **minADE** calculates the average  $\ell^2$  distance between predicted trajectory  $\hat{N}$  and ground truth over all time steps. **Miss Rate** reflects the proportion of scenarios where the

distance between the ground truth endpoint and the best-predicted endpoint exceeds 2.0 meters.

##### B. Quantitative Results

TABLE III: Quantitative Results I - This table presents the performance of QCNet using temporal ensembling with learning-based aggregation on the Argoverse 2 validation set. The baseline observes 50 past frames to predict 60 future steps (evaluated at intervals  $[10, 60]$ ). For a fair comparison, single-frame predictions at  $t = 10$  are made to match temporal ensembling by including the most recent agent states.

Method	minADE	minFDE	MissRate
QCNet [10]	0.50	0.99	10.73%
TempEns w/ LearnAgg + QCNet [10]	<b>0.48</b> (4.0%↓)	<b>0.94</b> (5.0%↓)	<b>9.57%</b> (1.16%↓)

We compare our proposed pipeline, Temporal Ensembling, with the strongest baseline, QCNet [10], on the Argoverse 2 dataset [14]. The results in Table I reveal that trajectory-level aggregations using KMeans, NMS, and TopK exhibit performance inferior to QCNet [10], contrary to our initial motivation. However, the temporal ensembling with learning-based aggregation shown in Table III surpasses all trajectory-level aggregation methods and even outperforms QCNet [10] in three metrics. Specifically, it achieves a 4% improvement in minADE, a 5% improvement in minFDE, and a 1.16% improvement in Miss Rate. These results align with our motivation, highlighting the feasibility of integrating predictions across multiple frames.

TABLE IV: Quantitative Results of Different Base Model Configurations - We investigated the effect of base model settings on performance. Both models were trained on a GeForce RTX 4090, differing only in the observation window and prediction horizon. The (40/70) base model, observing 40 frames and predicting 70 (evaluated at intervals  $[0, 60]$ ), enabled further temporal ensembling.

Method	minADE	minFDE	MissRate
QCNet [10] - Naive (50/60)	<b>0.65</b>	<b>1.27</b>	<b>0.16</b>
QCNet [10] - (40/70)	0.68	1.32	0.17

TABLE V: Quantitative Results II - The table compares our method's performance on the Argoverse 2 test set to SOTA approaches. For the official benchmark, evaluation is conducted at the  $[0, 60]$  interval. \* denotes the (40/70) base model. # indicates methods without publicly available code. All results exclude model ensembling for a fair comparison.

Method	minADE	minFDE	MissRate
GANet# (ICRA 2023) [30]	0.69	1.33	0.18
ProphNet# (CVPR 2023) [11]	0.65	1.31	0.17
QCNet* (CVPR 2023) [10]	0.68	1.32	0.17
TempEns w/ LearnAgg + QCNet* [10]	<b>0.65</b>	<b>1.27</b>	<b>0.16</b>

In the second quantitative result as shown in Table V, we evaluate our performance against the state-of-the-art methods on the Argoverse 2 test set [14]. There are two key findings. Firstly, reducing the observation window size can lead to a degradation in model performance (Table IV). Secondly, in subsequent experiments, incorporating our proposed approach to the base model with shorter observation periods demonstrates observable improvements across all evaluation metrics, reaffirming the effectiveness of our method.

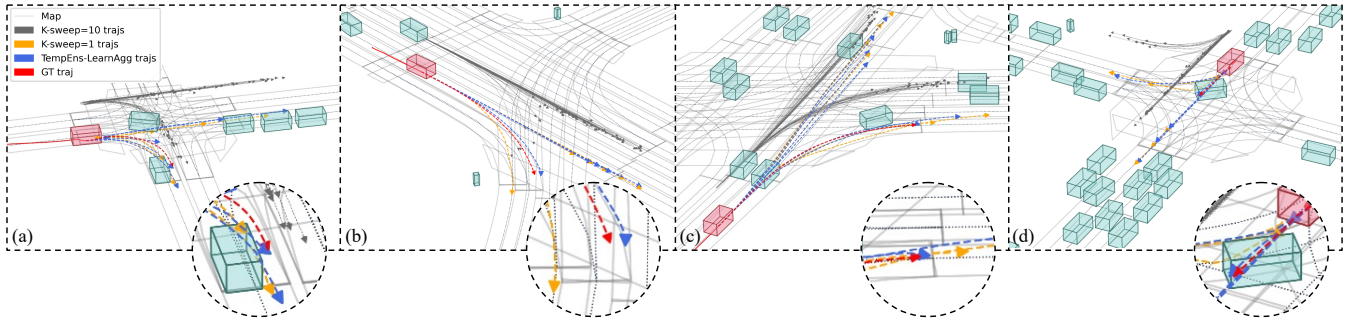


Fig. 6: Qualitative Results - Demonstrate the results of employing temporal ensembling with learning-based aggregation. The trajectories are color-coded: red for ground truth, orange for single-frame predictions, and blue for predictions using temporal ensembling with learning-based aggregation. Gray trajectories depict recent frame predictions, aiding in understanding the aggregation process.

### C. Analysis of Computational Overhead

TABLE VI: Computational Overhead - We measured inference time using a GeForce RTX 4090 in the **densest** traffic scene involving **190** agents and **169** map polygons.

Method	Online Inference Time
w/o TempEns	$80 \pm 1$ (ms)
w/ TempEns	$108 \pm 1$ (ms)

During training, our approach adopts a pre-trained QCNet [10] as the base model and only fine-tunes the newly added decoder. Practically, the resources needed for fine-tuning are significantly fewer compared to retraining a base model. During inference, temporal ensembling can be achieved with minimal computational resources, as it only requires a single forward pass. Historical mode queries are temporarily cached as the model progresses, enabling temporal ensembling by directly accessing recent data at the current time step. We also conducted experiments on inference time as shown in Table VI to demonstrate that a slight increase in computational resources leads to significant performance improvements.

### D. Qualitative Results

In Fig. 6, we showcase qualitative results from the Argoverse 2 validation set [14]. These cases highlight instances where the baseline method [10] (orange trajectories) failed, particularly in cases where the endpoint accuracy of the best trajectory exceeds a 2-meter error. Conversely, our approach (blue trajectories) successfully predicts the exact behaviors.

### E. Dual Ensembling with Varied Aggregation Techniques

TABLE VII: Exploring Model and Temporal Ensembling with Different Aggregation on the Argoverse 2 validation set. The baseline observes 50 past frames to predict 60 future steps, evaluated at intervals [10, 60).

Method	minADE	minFDE	MissRate
Baseline	0.50	0.99	10.73%
+ ModelEns w/ TrajAgg	0.50	0.98	10.20%
+ ModelEns w/ LearnAgg	0.49	0.96	9.84%
+ TempEns w/ LearnAgg	0.48	0.94	9.57%
+ ModelEns & TempEns w/ LearnAgg	<b>0.48</b>	<b>0.94</b>	<b>9.52%</b>

We explore the combined effects of model and temporal ensembling and their applicability when integrated. The

number of models in model ensembling is determined by computational resources. For our experiment, we train three model instances with varying epochs for model ensembling. In the initial model ensembling with trajectory-level aggregation experiment (Table VII's second row), predictions from three separate models are merged and combined using K-means, resulting in improvements in miss rate and minFDE, as expected. In the third row, we replaced the trajectory-level aggregation with our proposed learning-based aggregation. We observed that it outperformed trajectory-level aggregation, which we attribute to its effectiveness in integrating features from different models and then re-considering surrounding traffic information, thereby improving prediction accuracy. The fourth row of Table VII represents the results of the pipeline proposed in this paper. We conclude that both components are crucial. **Temporal Ensembling** contributes diversity, while **Learning-based Aggregation** enhances precision. In the final row of Table VII, we employ a dual approach by first averaging the mode queries  $q_{nt}^*$  across models during model ensembling and then applying the proposed temporal ensembling pipeline from nearby frames, which performs best across all three major metrics. Combining both pipelines yielded further performance enhancements.

### F. Alternative Base Model with DETR-like Architecture

TABLE VIII: Ablation Study on Alternative Base Models - This table showcases the performance of mmTrans using temporal ensembling with learning-based aggregation on the Argoverse 1 validation set. The base model uses 20 past frames to forecast 30 future steps, evaluated at intervals of [4, 30). To ensure a fair comparison, single-frame prediction at  $t = 4$  includes the most recent agent states to align with our approach.

Method	minADE	minFDE	MissRate
mmTrans [26]	0.62	0.90	7.38%
TempEns w/ LearnAgg + mmTrans	<b>0.61</b>	<b>0.88</b>	<b>6.82%</b>

We applied our proposed meta-algorithm, Temporal Ensembling with Learning-based Aggregation, to another well-renowned DETR-like model, mmTrans [13], to assess its impact on performance. Our results showed significant improvements in all three major performance indicators, further validating the effectiveness of our approach.

## V. CONCLUSION AND FUTURE WORKS

In this paper, we introduce a simple yet effective meta-algorithm, Temporal Ensembling with Learning-based Aggregation, to address the challenge of missing behaviors, effectively compensating single-frame predictions from multiple frames. Experimental results validate the effectiveness of the proposed approach, aiming to spark interest in exploring motion forecasting within a realistic streaming setting. A limitation of our method is its dependence on the assumption that nearby frames contain accurate predictions, making it less effective when predictions consecutively fail. For future work, we aim to comprehensively address this issue affecting safety by developing solutions that ensure robustness and accuracy in motion prediction across various scenarios and over time.

## REFERENCES

- [1] D. A. Pomerleau, "Alvin: An autonomous land vehicle in a neural network," *Advances in neural information processing systems*, vol. 1, 1988.
- [2] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *ROBOMECH journal*, vol. 1, no. 1, pp. 1–14, 2014.
- [3] W. Zhan, A. de La Fortelle, Y.-T. Chen, C.-Y. Chan, and M. Tomizuka, "Probabilistic prediction from planning perspective: Problem formulation, representation simplification and evaluation metric," in *2018 IEEE intelligent vehicles symposium (IV)*. IEEE, 2018, pp. 1150–1156.
- [4] J. Barceló *et al.*, *Fundamentals of traffic simulation*. Springer, 2010, vol. 145.
- [5] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, "Vectormet: Encoding hd maps and agent dynamics from vectorized representation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 525–11 533.
- [6] M. Liang, B. Yang, R. Hu, Y. Chen, R. Liao, S. Feng, and R. Urta-sun, "Learning lane graph representations for motion forecasting," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer, 2020, pp. 541–556.
- [7] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal *et al.*, "Scene transformer: A unified architecture for predicting multiple agent trajectories," *arXiv preprint arXiv:2106.08417*, 2021.
- [8] B. Varadarajan, A. Hefny, A. Srivastava, K. S. Refaat, N. Nayakanti, A. Cornman, K. Chen, B. Douillard, C. P. Lam, D. Anguelov *et al.*, "Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 7814–7821.
- [9] Z. Zhou, L. Ye, J. Wang, K. Wu, and K. Lu, "Hivt: Hierarchical vector transformer for multi-agent motion prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8823–8833.
- [10] Z. Zhou, J. Wang, Y.-H. Li, and Y.-K. Huang, "Query-centric trajectory prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17 863–17 873.
- [11] X. Wang, T. Su, F. Da, and X. Yang, "Prophet: Efficient agent-centric motion forecasting with anchor-informed proposals," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 995–22 003.
- [12] S. Shi, L. Jiang, D. Dai, and B. Schiele, "Motion transformer with global intention localization and local movement refinement," *Advances in Neural Information Processing Systems*, vol. 35, pp. 6531–6543, 2022.
- [13] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *European conference on computer vision*. Springer, 2020, pp. 213–229.
- [14] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, and J. Hays, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. [Online]. Available: <https://openreview.net/forum?id=vKQGe36av4k>
- [15] A. Barth and U. Franke, "Where will the oncoming vehicle be the next second?" in *2008 IEEE Intelligent Vehicles Symposium*. IEEE, 2008, pp. 1068–1073.
- [16] L. Chen, Y. Li, C. Huang, B. Li, Y. Xing, D. Tian, L. Li, Z. Hu, X. Na, Z. Li *et al.*, "Milestones in autonomous driving and intelligent vehicles: Survey of surveys," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 2, pp. 1046–1056, 2023.
- [17] H. Cui, V. Radosavljevic, F.-C. Chou, T.-H. Lin, T. Nguyen, T.-K. Huang, J. Schneider, and N. Djuric, "Multimodal trajectory predictions for autonomous driving using deep convolutional networks," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2090–2096.
- [18] B. Kim, S. H. Park, S. Lee, E. Khoshimjonov, D. Kum, J. Kim, J. S. Kim, and J. W. Choi, "Lapred: Lane-aware prediction of multi-modal future trajectories of dynamic agents," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 636–14 645.
- [19] Z. Huang, X. Mo, and C. Lv, "Multi-modal motion prediction with transformer-based neural network for autonomous driving," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2605–2611.
- [20] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid *et al.*, "Tnt: Target-driven trajectory prediction," in *Conference on Robot Learning*. PMLR, 2021, pp. 895–904.
- [21] J. Gu, C. Sun, and H. Zhao, "Densetnt: End-to-end trajectory prediction from dense goal sets," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 303–15 312.
- [22] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 1468–1476.
- [23] M. A. Ganaie, M. Hu, A. Malik, M. Tanveer, and P. Suganthan, "Ensemble deep learning: A review," *Engineering Applications of Artificial Intelligence*, vol. 115, p. 105151, 2022.
- [24] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [25] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*. Springer, 2020, pp. 683–700.
- [26] Y. Liu, J. Zhang, L. Fang, Q. Jiang, and B. Zhou, "Multimodal motion prediction with stacked transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7577–7586.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [28] S. Lee, S. Purushwalkam Shiva Prakash, M. Cogswell, V. Ranjan, D. Crandall, and D. Batra, "Stochastic multiple choice learning for training diverse deep ensembles," *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [29] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [30] M. Wang, X. Zhu, C. Yu, W. Li, Y. Ma, R. Jin, X. Ren, D. Ren, M. Wang, and W. Yang, "Ganet: Goal area network for motion forecasting," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1609–1615.