

VT-Former: An Exploratory Study on Vehicle Trajectory Prediction for Highway Surveillance through Graph Isomorphism and Transformer

Armin Danesh Pazho

University of North Carolina Charlotte
adaneshp@uncc.edu

Vinit Katariya

University of North Carolina Charlotte
vkatariy@uncc.edu

Ghazal Alinezhad Noghre

University of North Carolina Charlotte
galinezh@uncc.edu

Hamed Tabkhi

University of North Carolina Charlotte
htabkhiv@uncc.edu

Abstract

Enhancing roadway safety has become an essential computer vision focus area for Intelligent Transportation Systems (ITS). As a part of ITS, Vehicle Trajectory Prediction (VTP) aims to forecast a vehicle's future positions based on its past and current movements. VTP is a pivotal element for road safety, aiding in applications such as traffic management, accident prevention, work-zone safety, and energy optimization. While most works in this field focus on autonomous driving, with the growing number of surveillance cameras, another sub-field emerges for surveillance VTP with its own set of challenges. In this paper, we introduce VT-Former, a novel transformer-based VTP approach for highway safety and surveillance. In addition to utilizing transformers to capture long-range temporal patterns, a new Graph Attentive Tokenization (GAT) module has been proposed to capture intricate social interactions among vehicles. This study seeks to explore both the advantages and the limitations inherent in combining transformer architecture with graphs for VTP. Our investigation, conducted across three benchmark datasets from diverse surveillance viewpoints, showcases the State-of-the-Art (SotA) or comparable performance of VT-Former in predicting vehicle trajectories. This study underscores the potential of VT-Former and its architecture, opening new avenues for future research and exploration.

1. Introduction

Intelligent Transportation Systems (ITS) have emerged as a significant area of interest in research, with a notable classification distinguishing between autonomous driving [15, 19, 20, 22, 24, 42, 47, 62, 72, 73, 88] and surveillance applications [18, 40, 61, 65, 81, 90]. The subject of Ve-




Figure 1. A conceptual overview of the VT-Former. The input sequence, characterized by an observation horizon denoted as T_{OH} , undergoes tokenization via a novel graph-based tokenizer. The resultant tokens are subsequently processed through a transformer-based predictive model to forecast future trajectories.

hicle Trajectory Prediction (VTP), as a subset of ITS, has garnered significant attention in both areas. VTP is utilized within various higher-level applications among which are traffic management and control [36, 48, 52, 68, 80], accident and collision avoidance [25, 34, 54], workzone safety [36, 83, 86], route planning [67, 87], and vehicle anomaly detection [27, 33, 35, 64], etc.

The rapid advancements in computer vision and the widespread availability of surveillance cameras have opened up promising avenues and introduced novel techniques for advancing Surveillance Vehicle Trajectory Prediction (SVTP) methods [14, 28, 43]. This study focuses on SVTP in which the historical vehicle movement is used to forecast the future path, position, and behavior of a vehicle over a specific time horizon. While many studies focus on trajectory prediction from the perspective of autonomous driving and autonomous vehicles [5, 26, 56, 76], SVTP models approach it from surveillance angles, including the bird's eye view, high-angle, and eye-level perspectives [3, 36]. Another notable distinction between the two domains is that models in autonomous driving often pre-

dict multiple trajectories and select the most suitable one for evaluation [14, 58]. In contrast, SVTP methods typically rely on deterministic models that forecast a singular trajectory [3, 36, 46], which is more aligned with the practical requirements of highway safety scenarios.

There are several challenges and requirements associated with SVTP. Many applications such as traffic management or vehicle anomaly detection, require a balance between efficiency and accuracy [3, 44, 58]. This balance is pivotal as it directly impacts the real-world applicability and effectiveness of SVTP models in dynamic and often unpredictable environments. Another challenge is that vehicles interact with each other. Predicting how these interactions will unfold is complex, especially when combined with the uncertain dynamics of vehicle movements. Furthermore, a significant challenge stems from the design of most SVTP approaches for a bird's-eye view, whereas highway surveillance typically employs high-angle view cameras instead of bird's-eye views. Therefore, it is crucial to assess such models from the perspective of these more common real-world viewpoints as well.

In this paper, we present VT-Former, an innovative approach that leverages the power of transformers for vehicle trajectory prediction for highway surveillance applications. VT-Former also introduces Graph Attentive Tokenization (GAT) to capture the underlying social interactions among moving vehicles within the scene. The use of graphs has demonstrated significant success in pedestrian trajectory prediction, as evidenced by prior studies [3, 57]. This exploratory study aims to assess their effectiveness in the vehicle surveillance domain. Integrating these key components, GAT, and the decoder-only transformer, within VT-Former, offers a novel approach for SVTP. Fig. 1 shows a conceptual overview of VT-Former.

To demonstrate the benefits and limitations of VT-Former, we perform extensive experimentation using three major datasets: NGSIM [12, 13], CHD High-angle and CHD Eye-level [37] datasets, to comprehensively evaluate the performance of the VT-Former model from diverse perspectives. The evaluation encompassed three distinct angle variations corresponding to each dataset. Please note that VT-Former is designed for highway surveillance, not autonomous driving, making datasets like Argoverse [9, 84], nuScenes [8], PUP [30], and Waymo [21, 74] unrelated to this work. The outcomes underscore the advantages and potential, as well as the shortcomings of employing GAT alongside transformers for enhancing trajectory prediction accuracy and efficiency. On top of comparison with SotA models, we delved into a thorough analysis, comparing the effects of shortened historical trajectory data against extended historical horizons. This analysis explores the benefits and shortcomings of using a shortened observation horizon in cases where quick response time or lower prediction

latency is critical. Moreover, our ablation study demonstrates the advantages of utilizing a higher frame rate over the conventional setup, while preserving the existing ratio between the prediction horizon and the observed trajectory.

Overall, the contributions of this paper can be summarized as follows:

- Introducing VT-Former, exploring the combination of graph technology and transformer architecture as a new avenue and a novel approach for vehicle trajectory prediction for highway safety applications.
- Comprehensive assessments on benchmark datasets encompassing diverse surveillance camera angles, including bird's-eye view, high-angle view, and eye-level view.
- Exploring the influence of reduced observation horizons as well as higher frame rate on model accuracy, underscoring its effectiveness in scenarios demanding rapid trajectory prediction.

2. Related Works

In recent years, ITS have emerged as a focal point of extensive research endeavors within the scientific community [1, 23, 59, 69, 77]. Modern systems are increasingly integrating deep learning-based prediction mechanisms [3, 50, 63]. For many applications, path prediction is crucial as it helps in foreseeing future actions and locations beneficial for many downstream tasks. These predictive tools normally aim to operate in real-time [29, 36]. Such methodologies can be utilized for forecasting traffic densities [10, 31], improving the safety of work zones [70], and promoting autonomous vehicle technologies [71]. These predictive models also play a critical role in enabling advancements in smart monitoring and surveillance for enhanced security [66], improving pedestrian and worker safety in urban environments [57, 82], facilitating the identification and forecasting of anomalies within complex systems [55], and optimizing the design of transportation routes for increased efficiency and resource allocation [6].

Traditional methods, such as those based on Long Short-Term Memory (LSTM) networks, have been explored by Alahi et al. [2]. While effective, these methods often lack the flexibility and scalability of newer architectures. Graph Neural Networks (GNNs) have emerged as a powerful tool for trajectory prediction [17, 45], capitalizing on their ability to model relational data and complex interactions between entities. Specifically, Graph Convolutional Networks (GCN) [11, 71] allows for localized feature learning on graphs, making them apt for capturing spatial relationships between vehicles [51, 53]. The Graph Isomorphism Network (GIN) [57] further extends this capability by distinguishing different graph structures, enabling more precise trajectory predictions by considering nuanced vehicular interactions [41, 79, 89].

In recent studies, there has been a notable shift from

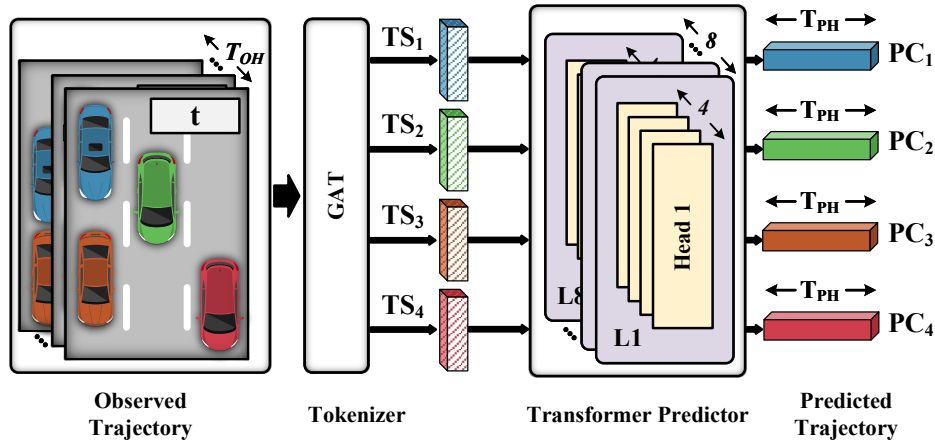


Figure 2. VT-Former at a glance: our approach begins with the application of Graph Attentive Tokenization (GAT) to create enriched token sequences that encapsulate social interactions. Subsequently, these enhanced sequences are fed through the Transformer Prediction module, which consists of 8 layers of transformer decoders, each equipped with 4 heads, generating predicted trajectories. VT-Former is an autoregressive sequence-to-sequence model that generates the output in multiple steps. However, for easier visualization, the output tokens are shown in the last step where they are completely generated. T_{OH} , TS_i , T_{PH} , and PC_i denote the observation horizon time, token sequence of the i^{th} vehicle, prediction horizon time, and the predicted trajectory, respectively.

Long Short-Term Memory (LSTM) and Convolutional Neural Networks (CNNs) to Transformers, as introduced by Vaswani et al. [78]. Transformers employ self-attention mechanisms for the analysis of sequential data making them suitable for many applications such as computer vision [39] and natural language processing [75]. Simultaneously, there have been endeavors to harness the power of Transformers in predicting trajectories within the realm of autonomous vehicles, enabling a better understanding of the temporal evolution of object movements [32, 91]. However, transformer-based methods remain largely unexplored within highway surveillance and safety applications. Moreover, their potential synergy with graphs to more effectively capture social nuances has yet to be investigated.

Despite the multitude of trajectory prediction algorithms proposed over the past decade, a limited number are explicitly designed for highway environments [4, 36]. Furthermore, the field of highway trajectory prediction is hampered by the scarcity of comprehensive datasets. Some recent works such as [37] have attempted to mitigate this and fill the gap. However, there remains a distinct requirement for specialized datasets tailored to highway-based applications. In general, there is a noticeable disparity between research endeavors and real-world applications in trajectory prediction especially when using already installed surveillance cameras. This underscores the potential for research and further development in this domain.

3. VT-Former Formulation and Methodology

At its core, VT-Former is a transformer-based approach for SVTP in roadways. Transformers are renowned for their exceptional generalization capabilities [78], enabling them to discern intricate long-term patterns within data and excel in tasks featuring extensive and diverse datasets. While the original Transformer model was computationally demanding and exhibited slower performance as it had both encoder and decoder components, numerous optimizations and model adaptations have been developed to enhance their efficiency [7, 38] by choosing only encoder or decoder as the core of their models. This has facilitated the utilization of transformers in applications necessitating low-latency responses. These attributes motivated our choice of the decoder-only transformer [7] as the core component of our algorithm.

VT-Former model utilizes a temporal window of the movement of a vehicle as input. This window is precisely determined by the spatial coordinates denoting the center of the vehicle’s bounding box and extends across the entirety of the observation horizon, creating a temporal window of the movements of the vehicle along the time axis:

$$C_i = [(x_i, y_i)^{t-T_{OH}+1}, (x_i, y_i)^{t-T_{OH}+2}, \dots, (x_i, y_i)^t] \quad (1)$$

where $(x_i, y_i)^t$ represents the center of the i^{th} vehicle’s bounding box and T_{OH} is the observation horizon. The model predicts the future location of vehicles through the prediction horizon (T_{PH}):

$$PC_i = [(x'_i, y'_i)^{t+1}, (x'_i, y'_i)^{t+2}, \dots, (x'_i, y'_i)^{t+T_{PH}}] \quad (2)$$

The overview of VT-Former can be seen in Fig. 2. Our proposed design consists of two main modules; the Graph Attentive Tokenization (GAT) module and the Transformer Prediction module. In the following subsections, we will delve into the details of the proposed architecture.

3.1. Graph Attention Tokenization

Graph Attention Tokenization (GAT) illustrated in Fig. 3 aims to extract enriched tokens that encapsulate the social interactions between the vehicles available in the scene. In order to accentuate the relative movements of the vehicles through the observation window we define the relative trajectory with respect to the first observed location:

$$\Delta C_i = [(x_i, y_i)^{t-T_{OH}+1} - (x_i, y_i)^{t-T_{OH}+1}, \\ (x_i, y_i)^{t-T_{OH}+2} - (x_i, y_i)^{t-T_{OH}+1}, \quad (3) \\ \dots, (x_i, y_i)^t - (x_i, y_i)^{t-T_{OH}+1}]$$

C_i and ΔC_i are concatenated to form the input of the fully connected layer which expands the dimension of the input features by a factor of 2. In the next step, we leverage Graph Isomorphism Network (GIN) [89] to capture the relationships between subjects available in the scene. GINs have proven highly effective in the context of processing graph-structured data, where nodes and edges signify entities and their relationships, respectively. In our formulation, each vehicle will be considered as a node in a fully connected graph. We chose to have a fully connected graph so that the Graph Network does not have any predefined biases and all relationships are learned by training. Another benefit of having a fully connected graph is that by performing just one step graph operation, the information will be propagated over all nodes and further steps are not necessary leading to a more efficient design. Inspired by [3], we design separate aggregation networks for the node feature and the neighbor features. The node feature aggregator is an MLP network with one hidden layer:

$$E'_v = Act(W_1 \cdot Act(W_0 \cdot (\beta \cdot E_v) + B_0) + B_1) \quad (4)$$

where E_v is the feature map for vehicle v , β is a learnable parameter, and W_0 , B_0 , W_1 and B_1 are the parameters of an MLP with one hidden layer. We have used the Leaky ReLu function as an activation function.

For aggregating neighbor features, we designed a more sophisticated network equipped with an attention mechanism:

$$E'_u = Act(W_2 \cdot E_u + B_2) \\ E''_u = A_s[A_C(E'_u) \times E'_u] \times [A_C(E'_u) \times E'_u] + E'_u \quad (5) \\ \text{where } u \in \mathcal{N}(v)$$

E_u is the feature map of neighbour vehicle, $\mathcal{N}(v)$ is the set of neighbour vehicles, W_2 and B_2 are the parameters of a fully connected layer, Act is Leaky ReLu function, and A_s and A_c are spatial attention and channel attention mechanism introduced by [85]. The attention mechanism helps the network to focus more on the features from neighbors that are most informative for predicting the future trajectory. The sum of E'_v and E''_u for all neighbor nodes is again concatenated with ΔC_i to emphasize relative movement constructing the predictor input sequence. In order to also add temporal ordering to the tokens in the predictor input sequence, we leverage temporal encoding proposed by [78] as depicted by Fig. 3. This process is necessary due to the fact that the Transformer Predictor module in the next step is not able to grasp the order of the sequence as suggested by [78].

3.2. Transformer Predictor Module

In the domain of time series data processing, Recurrent Neural Networks (RNNs) have traditionally been the primary choice. However, RNNs exhibit limitations related to their short-term memory, often constrained by challenges like vanishing gradients, which limits their ability to capture relationships between distant elements within a sequence. Several works also have tried using variations of Convolutional Neural Networks (CNNs) [3, 57] for finding temporal patterns, however, CNNs are not the preferred choice for modeling time series data because they lack the ability to effectively capture temporal dependencies, handle variable sequence lengths, and recognize long-range sequential patterns due to limited kernel size. Recently, a substantial paradigm shift has occurred with the emergence of transformers [78]. Transformers have garnered significant attention, predominantly due to their capability to capture long-range dependencies. Originally designed for Natural Language Processing (NLP) tasks, transformers have been adopted across a vast number of domains, including but not limited to computer vision, image and video processing, etc.

Hence, we have adopted a transformer-based architecture for more accurate vehicle trajectory prediction. The attention score inside the transformer layers is calculated using three learnable matrices Query (Q), Key (K), and Value (V) as described in [78]:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (6)$$

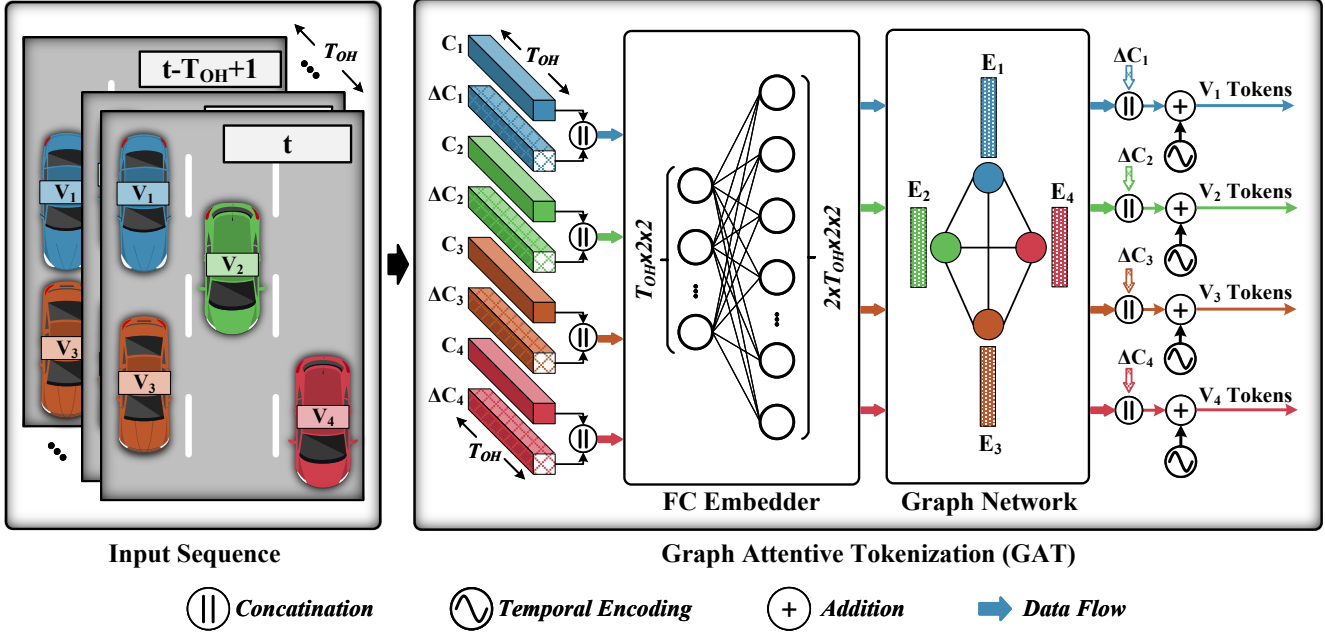


Figure 3. Graph Attentive Tokenization (GAT). GAT gets the past trajectory (C_i) of each vehicle (V_i) over the observation horizon (T_{OH}) for each vehicle and calculates the relative movement (ΔC_i). These two components are concatenated and passed through a fully connected layer, expanding the feature maps by a factor of 2. Subsequently, we employ a Graph Network to capture intricate interactions among vehicles within the scene. The output of this Graph Network is further concatenated with the relative movement, accentuating the temporal evolution within the sequence. Finally, to infuse temporal order into the token sequences, we leverage temporal encoding.

where d_k is the dimensionality of the Key matrix. Multiple attention heads are employed in transformers making them capable of attending to various input patterns. In our network, we have opted for a configuration with four attention heads. Following the attention mechanism, a Feed Forward Network (FFN) with a single hidden layer is utilized:

$$\text{FFN}(x) = \text{Act}(x \cdot W_3 + b_3) \cdot W_4 + b_4 \quad (7)$$

where W_3 , b_3 , W_4 , and b_4 are the parameters of the two fully connected layers.

The Transformer Predictor Module consists of 8 decoder layers stacked on top of each other. We chose the dimension of the FFN to be a size of 256. Within our design, as detailed in Sec. 3.1, we incorporated inter-vehicle attention in the tokenization process, although temporal attention among tokens was not considered. The multi-head attention mechanism in transformers is an ensemble of distinct attention mechanisms, with each head specializing in particular patterns. This adaptability empowers transformers to effectively capture a wide spectrum of data dependencies in the temporal dimension, ultimately improving the accuracy of predicted trajectories. We choose an autoregressive strategy where the model generates the sequence of predicted locations one step at a time in an ordered manner, where each prediction is influenced by the previous predictions leading

to more accurate prediction. To prevent the Transformer Predictor from peeking into future information, we employ future masking while training. The masking procedure is particularly important to ensure the causality of the system in parallel training mode.

4. Experimental Setup

All trainings have been conducted on a server equipped with 2x EPYC 7513 processors and 3x Nvidia A6000 GPUs with 256GB of memory.

4.1. Datasets

This study employs three datasets, each reflecting a distinct surveillance viewpoint (bird’s-eye, high-angle, and eye-level), to evaluate the performance and robustness of the VT-Former model.

4.1.1 NGSIM

The NGSIM dataset, specifically the US-101 [13] and I-80 [12] subsets, offer a rich source of real-world traffic scenarios and driver behaviors. As suggested by the name, this dataset is gathered from two distinct locations namely southbound US 101 in Los Angeles, CA, and eastbound I-80 in Emeryville, CA. With millions of data samples captured in a bird’s-eye view at 10 Hz, it provides a comprehen-

sive perspective of the surrounding environment. NGSIM data is originally collected from cameras positioned on buildings near the freeways, and it is then converted into a bird’s-eye format. The bounding box and tracking information are all hand-annotated and thus the noise is minimal in this dataset. This holds particular significance in practical, real-world scenarios, especially when dealing with pre-installed surveillance camera systems where manually annotated data is absent. Consequently, transitioning from training on hand-annotated data to operational deployment in the real world may result in a degradation in the model’s performance.

4.1.2 Carolina’s Highway Dataset

Carolina’s Highway Dataset (CHD) [37] is a comprehensive collection of vehicle trajectory data extracted from videos recorded in various locations across North Carolina and South Carolina, USA. It includes two distinct viewpoints: eye-level (CHD Eye-level) and high-angle (CHD High-angle), each included in a separate dataset. CHD not only provides trajectory information but also includes annotations for multi-vehicle detection and tracking. The data was recorded in full HD 60Hz at eight different locations, each chosen to capture a wide range of traffic patterns, vehicle behaviors, and lighting conditions. CHD offers a diverse perspective on incoming traffic, making it a valuable resource for understanding real-world highway scenarios. The annotations for CHD Eye-level and CHD High-angle datasets are derived through the application of deep learning models. Owing to the machine-based annotation process, these datasets exhibit a greater amount of noisy data in comparison to the NGSIM dataset. However, in contrast to NGSIM, these machine-generated annotations readily align with conditions and data obtainable in real-world settings.

4.2. Metrics

In alignment with previous studies, three established metrics, namely Average Displacement Error (ADE), Final Displacement Error (FDE), and Root Mean Squared Error (RMSE), are utilized to assess VT-Former’s performance.

4.2.1 Average Displacement Error (ADE)

ADE is a metric commonly used in the field of trajectory forecasting and motion prediction. It quantifies the accuracy of predicted trajectories for moving objects, such as vehicles or pedestrians by measuring the average Euclidean distance between the ground truth coordinates (x^t, y^t) and the predicted coordinates $(x^{t'}, y^{t'})$ across a defined prediction horizon (T_{PH}) considering all available N vehicles in the scene:

$$ADE = \frac{\sum_{i=1}^N \sum_{t=1}^{T_{PH}} \sqrt{(x^t - x^{t'})^2 + (y^t - y^{t'})^2}}{N \times T_{PH}} \quad (8)$$

In essence, it gauges how closely the predicted trajectory aligns with the actual trajectory, providing a numerical assessment of the model’s forecasting performance. A lower ADE value indicates more accurate predictions, while a higher value suggests less accurate or less precise trajectory forecasts. ADE is a fundamental measure in evaluating the quality and reliability of trajectory prediction models.

4.2.2 Final Displacement Error (FDE)

FDE quantifies the Euclidean distance between the ground truth final coordinates $(x^{T_{PH}}, y^{T_{PH}})$ and the predicted final coordinates $(x'^{T_{PH}}, y'^{T_{PH}})$ averaged over all available N vehicles:

$$FDE = \frac{\sum_{i=1}^N \sqrt{(x^{T_{PH}} - x'^{T_{PH}})^2 + (y^{T_{PH}} - y'^{T_{PH}})^2}}{N} \quad (9)$$

In essence, FDE assesses how accurately the model can forecast the precise destination or concluding point of an object’s trajectory. A smaller FDE value indicates more accurate predictions, signifying that the model successfully anticipates the final position, while a larger FDE value indicates less precise forecasting, implying a greater deviation from the actual endpoint.

4.2.3 Root Mean Squared Error (RMSE)

RMSE measures the square root of the average of the squared differences between the ground truth trajectory and the predicted trajectory at time step t over all available N vehicles:

$$RMSE^t = \sqrt{\frac{1}{N} \sum_{i=1}^N [(x^t - x^{t'})^2 + (y^t - y^{t'})^2]} \quad (10)$$

A smaller RMSE indicates that the model’s predictions are closer to the actual values, while a larger RMSE signifies a greater level of error in the predictions.

4.3. Training Strategy and Hyperparameters

In line with the conventions of previous research, our study utilizes an observation horizon of 15 timesteps and a prediction horizon of 25 timesteps (referred to as VT-Former_{LH}, where LH denotes Long History) across all datasets examined. Additionally, we define two more variations of VT-Former: the first version observes 10 timesteps and predicts 25 (VT-Former_{MH}, where MH stands for Medium

Table 1. Comparison with SotA methods (CS-LSTM [16], GRIP++ [46], STA-LSTM [49], DeepTrack [36], and Pishgu [3]) on NGSIM [12, 13]. VT-Former results are shown with different observation horizons of 15, 10, and 5 for VT-Former_{LH}, VT-Former_{MH}, and VT-Former_{SH} respectively. The results are in meters.

Model	ADE FDE		RMSE					Params (k)
			1s	2s	3s	4s	5s	
CS-LSTM	2.29	3.34	0.61	1.27	2.09	3.10	4.37	191
GRIP++	2.01	3.25	0.38	0.89	1.45	2.14	2.94	-
STA-LSTM	1.89	3.16	0.37	0.98	1.17	2.63	3.78	124
DeepTrack	2.01	3.21	0.47	1.08	1.83	2.75	3.89	109
VT-Former _{LH}	2.53	5.51	0.46	1.33	2.35	3.50	4.80	155
VT-Former _{MH}	2.24	5.15	0.37	1.13	2.10	3.26	4.62	141
VT-Former _{SH}	2.10	4.91	0.30	0.99	1.90	3.00	4.31	132

Table 2. Comparison with SotA methods on CHD [37] High-angle dataset. VT-Former results are shown with different observation horizons of 15, 10, and 5 for VT-Former_{LH}, VT-Former_{MH}, and VT-Former_{SH} respectively. The results are in pixels.

Model	ADE FDE		RMSE				
			1s	2s	3s	4s	5s
S-STGCNN [60]	31.87	98.46	9.74	21.83	29.01	42.34	82.14
GRIP++ [46]	36.32	100.89	3.40	6.67	14.32	28.02	123.04
Pishgu [3]	18.33	61.92	4.04	7.48	13.99	24.30	51.51
VT-Former _{LH}	25.95	87.21	7.60	17.35	22.90	27.97	66.39
VT-Former _{MH}	25.90	87.90	6.40	14.62	20.56	29.44	70.05
VT-Former _{SH}	25.33	88.99	5.67	12.96	19.12	29.83	70.72

History), and the second version observes 5 timesteps and predicts 25 (VT-Former_{SH}, where SH stands for Short History).

Given the disparate recording rates of these datasets, a standardization process is undertaken to achieve comparability with previous research, necessitating the down-sampling of both datasets to a common frame rate of 5 frames per second. For training the model for all datasets, Mean Squared Error (MSE) loss has been utilized.

Throughout all training sessions, the models undergo 80 training epochs, with a learning rate of 0.01, a weight decay of 0.0005, and a dropout rate of 0.2 consistently applied using Adam optimizer with batch size set to 16.

5. Results

5.1. Comparison with State-of-the-Art Methods

In this section, we conduct a comparative analysis between VT-Former and other trajectory prediction methods, leveraging three extensive highway datasets discussed in Sec. 4.1: NGSIM [12, 13], CHD High-angle [37] and CHD Eye-level [37]. Please note that VT-Former is only compared to deterministic models [3, 36, 46] and not models such as [58] that predict multiple future trajectories. Both CHD datasets were recently released, and as such, the ma-

Table 3. Comparison with SotA methods on CHD [37] Eye-level dataset. VT-Former results are shown with different observation horizons of 15, 10, and 5 for VT-Former_{LH}, VT-Former_{MH}, and VT-Former_{SH} respectively. The results are in pixels.

Model	ADE FDE		RMSE				
			1s	2s	3s	4s	5s
S-STGCNN [60]	24.33	95.22	4.32	9.15	15.93	29.05	68.32
GRIP++ [46]	44.27	129.58	4.42	12.86	24.31	35.04	145.17
Pishgu [3]	37.99	123.69	4.98	13.58	26.61	50.31	106.45
VT-Former _{LH}	34.88	100.59	6.71	17.24	28.70	45.86	82.00
VT-Former _{MH}	27.44	85.45	5.19	12.90	21.38	35.09	68.60
VT-Former _{SH}	21.86	66.28	5.42	12.69	18.81	26.67	53.05




Figure 4. Qualitative performance of VT-Former on CHD, featuring a scene with a merging lane from the left and a mild right curve. The angle between predicted, observed, and actual trajectories in samples A and B demonstrates VT-Former’s ability to predict the correct path with respect to road geometry. However, it struggles with the high acceleration of vehicle 2 in sample B. Samples are cropped for clarity.

majority of SotA models have not been evaluated on this specific dataset. To bridge this gap, we have used the available code repositories of leading SotA models, reconfiguring them for training and evaluation on the CHD datasets.

As discussed in Sec. 1, this exploratory study assesses the advantages and limitations of using transformers and graphs for SVTP. The results presented in Tab. 1, Tab. 2, and Tab. 3, along with the qualitative findings such as Fig. 4, collectively illustrate the potential of the VT-Former for highway trajectory prediction. Although the VT-Former did not achieve high results on the NGSIM dataset, it performed notably well as the second-best model on both the CHD High-angle and CHD Eye-level datasets. These findings indicate that while Graph Isomorphism Networks (GIN) and transformers are viable for SVTP, there is still potential for further refinement. More specifically, the data can lead to the hypothesis that the complexity of the graph significantly impacts the outcomes. Therefore, practices such as transitioning from fully connected graphs to more dynamic and optimized structures, along with implementing techniques such as neighborhood selection, could enhance performance.

5.2. Observation Horizon Analysis

In this section, we conduct a comprehensive investigation into the influence of different observation horizons to discern the optimal configuration for each dataset. The observation horizon is varied, including 5 timesteps for VT-

Former_{SH} (translating to 1 second for 5Hz), 10 timesteps for VT-Former_{MH} (translating to 2 second for 5Hz) 15 timesteps for VT-Former_{LH} (translating to 3 second for 5Hz). Subsequently, we train and assess the performance of VT-Former on the NGSIM [12, 13], CHD High-angle [37], and CHD Eye-level [37] datasets. The results are presented in Tab. 1, Tab. 2, and Tab. 3 for the respective datasets. This specific assessment is especially important for applications that require faster prediction, such as collision avoidance or anomaly detection.

The results on all three datasets, as depicted in Tab. 1, Tab. 2, and Tab. 3, offer insights into the benefits of short-term observation horizon. Here, it is evident that a shorter observation horizon is more advantageous. For example, on the CHD Eye-level dataset, when deploying VT-Former_{SH}, we observe a remarkable enhancement of 37.3% in FDE and 34.1% in ADE compared to VT-Former_{LH} results. This outcome underscores the practicality of employing shorter observation horizons in the eye-level setup. Similar trends are seen for both NGSIM and CHD High-angle datasets. This phenomenon can be attributed to the nature of trajectory prediction. In longer trajectories, maintaining a high temporal resolution (detail at very small time scales) might be challenging due to computational constraints or model architecture limits. This loss of detail can lead to poorer predictions as the model might miss out on subtle cues that are more apparent in shorter, high-resolution trajectories. This particularly shows that with a lower observation horizon, VT-Former is able to predict future trajectories better.

In summary, the empirical evidence presented in this study shows that, in practical, real-world scenarios, the adoption of shorter observation horizons is an effective strategy for decreasing the negative impact of chaotic driving behavior. This approach not only leads to more accurate trajectory prediction outcomes but also offers the advantage of expeditious prediction. By reducing the temporal extent of observation, the model can effectively focus on details at smaller scales, maintaining a high temporal resolution, thus enhancing the precision of its predictions and enabling it to respond more swiftly. This finding underscores the practical utility of optimizing observation horizons in trajectory prediction models, particularly in applications where the timely and precise forecasting of object movements is important.

6. Ablation Study

In this section, we perform ablation testing to evaluate and compare the effect of higher data rate on the performance of the model, including the Observation Horizon test same as Sec. 5.2. We utilize the NGSIM dataset as it offers a 10Hz sampling rate compared to conventional 5Hz sampling. The performance of VT-Former and Pishgu [3] with an observation horizon of 15 timesteps (translating to 1.5 seconds for 10Hz) and a prediction horizon of 25 timesteps (trans-

Table 4. Comparison of VT-Former and Pishgu on NGSIM dataset sampled at 10 Hz. VT-Former results are shown with different observation horizons of 15, 10, and 5 for VT-Former_{LH}, VT-Former_{MH}, and VT-Former_{SH} respectively. Pishgu results are from its updated repository. The results are in meters.

Model	RMSE						
	ADE	FDE	0.5s	1.0s	1.5s	2.0s	2.5s
Pishgu [3]	0.88	1.96	0.15	0.46	0.82	1.25	1.74
VT-Former _{LH}	0.81	1.80	0.14	0.42	0.76	1.16	1.61
VT-Former _{MH}	0.85	1.89	0.16	0.47	0.84	1.27	1.75
VT-Former _{SH}	0.79	1.78	0.13	0.41	0.76	1.16	1.62

lating to 2.5 seconds for 10Hz) on the NGSIM dataset is presented in Tab. 4. VT-Former shows a 7.95% improvement in ADE and 8.16% improvement in FDE over Pishgu. A similar trend is observed in RMSE with the most pronounced improvement of 7.47% from Pishgu’s 1.74 meters to 1.61 meters over a 2.5-second prediction horizon. This shows that at a higher data rate, VT-Former not only outperforms Pishgu but achieves higher accuracy with respect to RMSE at further timesteps.

Additionally, the impact of varying observation horizons, including 5 timesteps for VT-Former_{SH} (translating to 0.5 seconds for 10Hz), 10 timesteps for VT-Former_{MH} (translating to 1 second for 10Hz) 15 timesteps for VT-Former_{LH} (translating to 1.5 seconds for 10Hz), is assessed and detailed in Tab. 4. These results exhibit similar trends to those discussed in Sec. 5.2, where a shorter observation horizon correlates with improved performance in predicting future trajectories.

7. Conclusion

This work introduced VT-Former, an innovative trajectory prediction algorithm that leverages transformer architectures in conjunction with a novel tokenization scheme utilizing graph attention mechanisms denoted as Graph Attentive Tokenization (GAT). Extensive experimentation was conducted across three benchmark datasets illustrating the performance of the model, as well as assessing the model’s generalizability and robustness. Further experimentation has been conducted to evaluate VT-Former’s performance under shorter observed horizons. This study underscores the potential of combining graphs with transformers, opening a new path for future research in the field.

Acknowledgment

This research is supported by the National Science Foundation (NSF) under Award No. 1932524.

References

- [1] Mohammad Abdullah Al Faruque, Shih-Yuan Yu, Arnav Vaibhav Malawade, Deepan Muthirayan, and Pramod P Khargonekar. Methods for spatio-temporal scene-graph embedding for autonomous vehicle applications, 2023. US Patent App. 18/154,987. [2](#)
- [2] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016. [2](#)
- [3] Ghazal Alinezhad Noghre, Vinit Katariya, Armin Danesh Pazho, Christopher Neff, and Hamed Tabkhi. Pishgu: Universal path prediction network architecture for real-time cyber-physical edge systems. In *Proceedings of the ACM/IEEE 14th International Conference on Cyber-Physical Systems (with CPS-IoT Week 2023)*, pages 88–97, 2023. [1](#), [2](#), [4](#), [7](#), [8](#)
- [4] Florent Alché and Arnaud de La Fortelle. An lstm network for highway trajectory prediction. In *2017 IEEE 20th international conference on intelligent transportation systems (ITSC)*, pages 353–359. IEEE, 2017. [3](#)
- [5] Neel P Bhatt, Amir Khajepour, and Ehsan Hashemi. Mpc-pf: Socially and spatially aware object trajectory prediction for autonomous driving systems using potential fields. *IEEE Transactions on Intelligent Transportation Systems*, 2023. [1](#)
- [6] Justin M Bradley and Ella M Atkins. Optimization and control of cyber-physical vehicle systems. *Sensors*, 15(9):23020–23049, 2015. [2](#)
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [3](#)
- [8] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [9] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#)
- [10] Dawei Chen. Research on traffic flow prediction in the big data environment based on the improved rbf neural network. *IEEE Transactions on Industrial Informatics*, 13(4):2000–2008, 2017. [2](#)
- [11] Hao Cheng, Mengmeng Liu, Lin Chen, Hellward Broszio, Monika Sester, and Michael Ying Yang. Gatraj: A graph and attention-based multi-agent trajectory prediction model. *ISPRS Journal of Photogrammetry and Remote Sensing*, 205:163–175, 2023. [2](#)
- [12] James Colyar and John Halkias. Next generation simulation (NGSIM), Interstate 80 freeway dataset. FHWA-HRT-06-137, 2006. [2](#), [5](#), [7](#), [8](#)
- [13] James Colyar and John Halkias. Next generation simulation (NGSIM), US Highway-101 dataset. FHWA-HRT-07-030., 2007. [2](#), [5](#), [7](#), [8](#)
- [14] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2090–2096. IEEE, 2019. [1](#), [2](#)
- [15] Henggang Cui, Fang-Chieh Chou, Jake Charland, Carlos Vallespi-Gonzalez, and Nemanja Djuric. Uncertainty-aware estimation of vehicle orientation for self-driving applications. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 2660–2666. IEEE, 2021. [1](#)
- [16] Nachiket Deo and Mohan M Trivedi. Convolutional social pooling for vehicle trajectory prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 1468–1476, 2018. [7](#)
- [17] Nachiket Deo, Eric Wolff, and Oscar Beijbom. Multimodal trajectory prediction conditioned on lane-graph traversals. In *Conference on Robot Learning*, pages 203–212. PMLR, 2022. [2](#)
- [18] Esma Dilek and Murat Dener. Computer vision applications in intelligent transportation systems: a survey. *Sensors*, 23(6):2938, 2023. [1](#)
- [19] Nemanja Djuric, Vladan Radosavljevic, Henggang Cui, Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, and Jeff G. Schneider. Short-term motion prediction of traffic actors for autonomous driving using deep convolutional networks. 2018. [1](#)
- [20] Nemanja Djuric, Henggang Cui, Zhaoen Su, Shangxuan Wu, Huahua Wang, Fang-Chieh Chou, Luisa San Martin, Song Feng, Rui Hu, Yang Xu, et al. Multixnet: Multiclass multistage multimodal motion prediction. In *2021 IEEE Intelligent Vehicles Symposium (IV)*, pages 435–442. IEEE, 2021. [1](#)
- [21] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R. Qi, Yin Zhou, Zoey Yang, Aurélien Chouard, Pei Sun, Jiquan Ngiam, Vijay Vasudevan, Alexander McCauley, Jonathon Shlens, and Dragomir Anguelov. 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)*, pages 9710–9719, 2021. [2](#)
- [22] Sudeep Fadadu, Shreyash Pandey, Darshan Hegde, Yi Shi, Fang-Chieh Chou, Nemanja Djuric, and Carlos Vallespi-Gonzalez. Multi-view fusion of sensor data for improved perception and prediction in autonomous driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2349–2357, 2022. [1](#)
- [23] Georgios Fainekos, Hani Ben Amor, Adel Dokhanchi, and Jyotirmoy Deshmukh. Systems and methods for evaluating perception systems for autonomous vehicles using quality temporal logic, 2023. US Patent 11,586,914. [2](#)

- [24] Shivam Gautam, Gregory P Meyer, Carlos Vallespi-Gonzalez, and Brian C Becker. Sdvtracker: Real-time multi-sensor association and tracking for self-driving vehicles. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3012–3021, 2021. **1**
- [25] Adam Houenou, Philippe Bonnifait, Véronique Cherfaoui, and Wen Yao. Vehicle trajectory prediction based on motion model and maneuver recognition. In *2013 IEEE/RSJ international conference on intelligent robots and systems*, pages 4363–4369. IEEE, 2013. **1**
- [26] Chih-Chung Hsu, Li-Wei Kang, Shih-Yu Chen, I-Shan Wang, Ching-Hao Hong, and Chuan-Yu Chang. Deep learning-based vehicle trajectory prediction based on generative adversarial network for autonomous driving applications. *Multimedia Tools and Applications*, 82(7):10763–10780, 2023. **1**
- [27] Shihong Ed Huang, Yiheng Feng, and Henry X Liu. A data-driven method for falsified vehicle trajectory identification by anomaly detection. *Transportation research part C: emerging technologies*, 128:103196, 2021. **1**
- [28] Yanjun Huang, Jiatong Du, Ziru Yang, Zewei Zhou, Lin Zhang, and Hong Chen. A survey on trajectory-prediction methods for autonomous driving. *IEEE Transactions on Intelligent Vehicles*, 7(3):652–674, 2022. **1**
- [29] Juan Isern, Francisco Barranco, Daniel Deniz, Juho Lesonen, Jari Hannuksela, and Richard R. Carrillo. Reconfigurable cyber-physical system for critical infrastructure protection in smart cities via smart video-surveillance. *Pattern Recognit. Lett.*, 140:303–309, 2020. **2**
- [30] Boris Ivanovic, Kuan-Hui Lee, Pavel Tokmakov, Blake Wulfe, Rowan Mcllister, Adrien Gaidon, and Marco Pavone. Heterogeneous-agent trajectory forecasting incorporating class uncertainty. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12196–12203. IEEE, 2022. **2**
- [31] Young-Seon Jeong, Young-Ji Byon, Manoel Mendonca Castro-Neto, and Said M. Easa. Supervised weighting-online learning algorithm for short-term traffic flow prediction. *IEEE Transactions on Intelligent Transportation Systems*, 14(4):1700–1707, 2013. **2**
- [32] Xiaosong Jia, Penghao Wu, Li Chen, Yu Liu, Hongyang Li, and Junchi Yan. Hdgt: Heterogeneous driving graph transformer for multi-agent trajectory prediction via scene encoding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):13860–13875, 2023. **3**
- [33] Ruochen Jiao, Juyang Bai, Xiangguo Liu, Takami Sato, Xiaowei Yuan, Qi Alfred Chen, and Qi Zhu. Learning representation for anomaly detection of vehicle trajectories. *arXiv preprint arXiv:2303.05000*, 2023. **1**
- [34] Chang Mook Kang, Soo Jung Jeon, Seung-Hi Lee, and Chung Choo Chung. Parametric trajectory prediction of surrounding vehicles. In *2017 IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, pages 26–31. IEEE, 2017. **1**
- [35] Gözde Boztepe Karataş, Pinar Karagoz, and Orhan Ayran. Trajectory pattern extraction and anomaly detection for maritime vessels. *Internet of Things*, 16:100436, 2021. **1**
- [36] Vinit Katariya, Mohammadreza Baharani, Nichole Morris, Omidreza Shoghli, and Hamed Tabkhi. Deeptrack: Lightweight deep learning for vehicle trajectory prediction in highways. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):18927–18936, 2022. **1, 2, 3, 7**
- [37] Vinit Katariya, Ghazal Alinezhad Noghre, Armin Danesh Pazho, and Hamed Tabkhi. Charlotte highway dataset. https://github.com/TeCSAR-UNCC/Carolinas_Dataset, 2023. Dataset. **2, 3, 6, 7, 8**
- [38] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, page 2, 2019. **3**
- [39] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM Comput. Surv.*, 54(10s), 2022. **3**
- [40] Gahyun Kim, Ho Gi Jung, and Jae Kyu Suhr. Cnn-based vehicle bottom face quadrilateral detection using surveillance cameras for intelligent transportation systems. *Sensors*, 23(15):6688, 2023. **1**
- [41] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. **2**
- [42] Ankit Laddha, Shivam Gautam, Stefan Palombo, Shreyash Pandey, and Carlos Vallespi-Gonzalez. Mvfusenet: Improving end-to-end object detection and motion forecasting through multi-view fusion of lidar data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2865–2874, 2021. **1**
- [43] Florin Leon and Marius Gavrilescu. A review of tracking and trajectory prediction methods for autonomous driving. *Mathematics*, 9(6):660, 2021. **1**
- [44] Ruonan Li, Yang Qin, Jingbo Wang, and Hongye Wang. Amgb: Trajectory prediction using attention-based mechanism gcn-bilstm in iov. *Pattern Recognition Letters*, 169:17–27, 2023. **2**
- [45] Xin Li, Xiaowen Ying, and Mooi Choo Chuah. Grip: Graph-based interaction-aware trajectory prediction. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 3960–3966. IEEE, 2019. **2**
- [46] Xin Li, Xiaowen Ying, and Mooi Choo Chuah. Grip++: Enhanced graph-based interaction-aware trajectory prediction for autonomous driving. *arXiv preprint arXiv:1907.07792*, 2019. **2, 7**
- [47] Junwei Liang, He Zhu, Enwei Zhang, and Jun Zhang. Stargazer: A transformer-based driver action detection system for intelligent transportation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3160–3167, 2022. **1**
- [48] Yuebing Liang and Zhan Zhao. Nettraj: A network-based vehicle trajectory prediction model with directional representation and spatiotemporal attention mechanisms. *IEEE Transactions on Intelligent Transportation Systems*, 23(9):14470–14481, 2021. **1**

- [49] Lei Lin, Weizi Li, Huikun Bi, and Lingqiao Qin. Vehicle trajectory prediction using lstms with spatial-temporal attention mechanisms. *IEEE Intelligent Transportation Systems Magazine*, 14(2):197–208, 2021. 7
- [50] Bufan Liu, Yingfeng Zhang, Geng Zhang, and Pai Zheng. Edge-cloud orchestration driven industrial smart product-service systems solution design based on cps and iiot. *Advanced Engineering Informatics*, 42:100984, 2019. 2
- [51] Yongkang Liu, Xuewei Qi, Emrah Akin Sisbot, and Kentaro Oguchi. Multi-agent trajectory prediction with graph attention isomorphism neural network. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, pages 273–279. IEEE, 2022. 2
- [52] Gongyuan Lu, Yu Marco Nie, Xiaobo Liu, and Denghui Li. Trajectory-based traffic management inside an autonomous vehicle zone. *Transportation Research Part B: Methodological*, 120:76–98, 2019. 1
- [53] Pei Lv, Wentong Wang, Yunxin Wang, Yuzhen Zhang, Mingliang Xu, and Changsheng Xu. Ssagcn: social soft attention graph convolution network for pedestrian trajectory prediction. *IEEE transactions on neural networks and learning systems*, 2023. 2
- [54] Nengchao Lyu, Jiaqiang Wen, Zhicheng Duan, and Chaozhong Wu. Vehicle trajectory prediction and cut-in collision warning model in a connected vehicle environment. *IEEE Transactions on Intelligent Transportation Systems*, 23(2):966–981, 2020. 1
- [55] Amir Markovitz, Gilad Sharir, Itamar Friedman, Lih Zelnik-Manor, and Shai Avidan. Graph embedded pose clustering for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10539–10547, 2020. 2
- [56] Rowan McAllister, Blake Wulfe, Jean Mercat, Logan Ellis, Sergey Levine, and Adrien Gaidon. Control-aware prediction objectives for autonomous driving. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 01–08. IEEE, 2022. 1
- [57] Matías Mendieta and Hamed Tabkhi. Carpe posterum: A convolutional approach for real-time pedestrian path prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2346–2354, 2021. 2, 4
- [58] Kaouther Messaoud, Itheri Yahiaoui, Anne Verroust-Blondet, and Fawzi Nashashibi. Attention based vehicle trajectory prediction. *IEEE Transactions on Intelligent Vehicles*, 6(1):175–185, 2020. 2, 7
- [59] Wilbur Michael, Amutheezan Sivagnanam, Afiya Ayman, Samitha Samaranayeke, Abhishek Dubey, and Aron Laszka. Artificial intelligence for smart transportation. *Vorobeychik, Yevgeniy., and Mukhopadhyay, Ayan.,(Eds.).(2023). Artificial Intelligence and Society. ACM Press Subjects: Artificial Intelligence (cs. AI) Cite as: arXiv: 2308.07457 [cs. AI](or arXiv: 2308.07457 v1 [cs. AI] for this version) https://doi.org/10.48550/arXiv.2308.07457 Focus to learn more, 2023. 2*
- [60] Abdulllah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14424–14432, 2020. 7
- [61] Akhmad Yusuf Nasirudin, Koredianto Usman, and Suryo Adhi Wibowo. Surveillance system scheme using multi-detection attribute with optimized neural network algorithm on intelligent transportation system. *eProceedings of Engineering*, 10(3), 2023. 1
- [62] Pha Nguyen, Kha Gia Quach, Chi Nhan Duong, Ngan Le, Xuan-Bac Nguyen, and Khoa Luu. Multi-camera multiple 3d object tracking on the move for autonomous vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2569–2578, 2022. 1
- [63] Felix O Olowononi, Danda B Rawat, and Chunmei Liu. Resilient machine learning for networked cyber physical systems: A survey for machine learning security to securing machine learning for cps. *IEEE Communications Surveys & Tutorials*, 23(1):524–552, 2020. 2
- [64] Billy Peralta, Richard Soria, Orietta Nicolis, Fabrizio Ruggeri, Luis Caro, and Andrés Bronfman. Outlier vehicle trajectory detection using deep autoencoders in santiago, chile. *Sensors*, 23(3):1440, 2023. 1
- [65] Arman Syah Putra, Harco Leslie Hendric Spits Warnars, Ford Lumban Gaol, Benfano Soewito, and Edi Abdurachman. A proposed surveillance model in an intelligent transportation system (its). In *2018 Indonesian association for pattern recognition international conference (INAPR)*, pages 156–160. IEEE, 2018. 1
- [66] Lakshmana Kumar Ramasamy, Firoz Khan, Mohammad Shah, Balusupati Veera Venkata Siva Prasad, Celestine Iwendi, and Cresantus Biamba. Secure smart wearable computing through artificial intelligence-enabled internet of things and cyber-physical systems for health monitoring. *Sensors*, 22(3):1076, 2022. 2
- [67] Punit Rathore, Dheeraj Kumar, Sutharshan Rajasegarar, Marimuthu Palaniswami, and James C Bezdek. A scalable framework for trajectory prediction. *IEEE Transactions on Intelligent Transportation Systems*, 20(10):3860–3874, 2019. 1
- [68] Luca Rossi, Andrea Ajmar, Marina Paolanti, and Roberto Pierdicca. Vehicle trajectory prediction and generation using lstm models and gans. *Plos one*, 16(7):e0253868, 2021. 1
- [69] Giulio Rossolini, Federico Nesti, Gianluca D’Amico, Saasha Nair, Alessandro Biondi, and Giorgio Buttazzo. On the real-world adversarial robustness of real-time semantic segmentation models for autonomous driving. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2023. 2
- [70] Sepehr Sabeti, Omidreza Shoghli, Mohammadreza Baharani, and Hamed Tabkhi. Toward ai-enabled augmented reality to enhance the safety of highway work zones: Feasibility, requirements, and challenges. *Advanced Engineering Informatics*, 50:101429, 2021. 2
- [71] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco 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*, pages 683–700. Springer, 2020. 2
- [72] Apoorv Singh. Surround-view vision-based 3d detection for autonomous driving: A survey. In *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 3235–3244. IEEE, 2023. 1
- [73] Apoorv Singh. Transformer-based sensor fusion for autonomous driving: A survey. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3312–3317, 2023. 1
- [74] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [75] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014. 3
- [76] Luqi Tang, Fuwu Yan, Bin Zou, Wenbo Li, Chen Lv, and Kewei Wang. Trajectory prediction for autonomous driving based on multiscale spatial-temporal graph. *IET Intelligent Transport Systems*, 17(2):386–399, 2023. 1
- [77] Francis Tiausas, Keiichi Yasumoto, Jose Paolo Talusan, Hayato Yamana, Hirozumi Yamaguchi, Shameek Bhattacharjee, Abhishek Dubey, and Sajal K Das. Hprop: Hierarchical privacy-preserving route planning for smart cities. *ACM Transactions on Cyber-Physical Systems*, 2023. 2
- [78] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3, 4
- [79] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. 2
- [80] Eason Wang, Henggang Cui, Sai Yalamanchi, Mohana Moorthy, and Nemanja Djuric. Improving movement predictions of traffic actors in bird’s-eye view models using gans and differentiable trajectory rasterization. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2340–2348, 2020. 1
- [81] Qian Wang, Cai Guo, Hong-Ning Dai, and Min Xia. Variant-depth neural networks for deblurring traffic images in intelligent transportation systems. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 1
- [82] Shuai Wang, Xiaoyu Li, Wei Chen, Weiqiang Fan, and Zijian Tian. An intelligent vision-based method of worker identification for industrial internet of things (iiot). *Wireless Communications and Mobile Computing*, 2022, 2022. 2
- [83] Jinxian Weng, Qiang Meng, and Xuedong Yan. Analysis of work zone rear-end crash risk for different vehicle-following patterns. *Accident Analysis & Prevention*, 72:449–457, 2014. 1
- [84] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Curran, 2021. 2
- [85] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 4
- [86] Renfei Wu, Xunjia Zheng, Yongneng Xu, Wei Wu, Guopeng Li, Qing Xu, and Zhuming Nie. Modified driving safety field based on trajectory prediction model for pedestrian–vehicle collision. *Sustainability*, 11(22):6254, 2019. 1
- [87] Zhuoqun Xia, Zhenzhen Hu, and Junpeng Luo. Uptp vehicle trajectory prediction based on user preference under complexity environment. *Wireless Personal Communications*, 97:4651–4665, 2017. 1
- [88] Cuichun Xu, Peter Kingston, and Nemanja Djuric. Improved tracking of articulated vehicles in self-driving applications using phantom observations. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pages 3699–3705. IEEE, 2023. 1
- [89] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018. 2, 4
- [90] Hao Frank Yang, Jiarui Cai, Chenxi Liu, Ruimin Ke, and Yinhai Wang. Cooperative multi-camera vehicle tracking and traffic surveillance with edge artificial intelligence and representation learning. *Transportation research part C: emerging technologies*, 148:103982, 2023. 1
- [91] Kunpeng Zhang, Xiaoliang Feng, Lan Wu, and Zhengbing He. Trajectory prediction for autonomous driving using spatial-temporal graph attention transformer. *IEEE Transactions on Intelligent Transportation Systems*, 23(11):22343–22353, 2022. 3