







High-Order Evolving Graphs for Enhanced Representation of Traffic Dynamics

Aditya Humnabadkar¹, Arindam Sikdar², Benjamin Cave³, Huaizhong Zhang, Paul Bakaki, and Ardhendu Behera¹

Department of Computer Sciences, Edge Hill University, Ormskirk, United Kingdom

Abstract. We present an innovative framework for traffic dynamics analysis using High-Order Evolving Graphs, designed to improve spatio-temporal representations in autonomous driving contexts. Our approach constructs temporal bidirectional bipartite graphs that effectively model the complex interactions within traffic scenes in real-time. By integrating Graph Neural Networks (GNNs) with high-order multi-aggregation strategies, we significantly enhance the modeling of traffic scene dynamics, providing a more accurate and detailed analysis of these interactions. Additionally, we incorporate inductive learning techniques inspired by the GraphSAGE framework, enabling our model to adapt to new and unseen traffic scenarios without the need for retraining, thus ensuring robust generalization. Through extensive experiments on the ROAD and ROAD Waymo datasets, we establish a comprehensive baseline for further developments, demonstrating the potential of our method in accurately capturing traffic behavior. Our results emphasize the value of high-order statistical moments and feature-gated attention mechanisms in improving traffic behavior analysis, laying the groundwork for advancing autonomous driving technologies. Our source code will be made available at: https://github.com/Addy-1998/High_Order_Graphs

Keywords: traffic dynamics representation · high-order evolving graphs · graph neural networks · multi-aggregation · temporal bidirectional bipartite graphs

1 Introduction

Recognizing driving activities and actions from video and image data is a crucial task in computer vision and pattern recognition, particularly for enhancing the safety of autonomous vehicles and preventing collisions. This technology is integral to Advanced Driver Assistance Systems (ADAS), contributing to the development of features such as adaptive cruise control and lane-keeping assistance, which are essential for maintaining safe vehicle distances and providing timely alerts for lane changes [38]. The complexity of driving scene-based activity detection arises from the highly dynamic and unpredictable nature of traffic environments, especially in urban areas where the overall scene dynamics, including the movement and interaction of various road users such as vehicles, pedestrians, and cyclists, play a critical role [5, 7]. These scene dynamics are

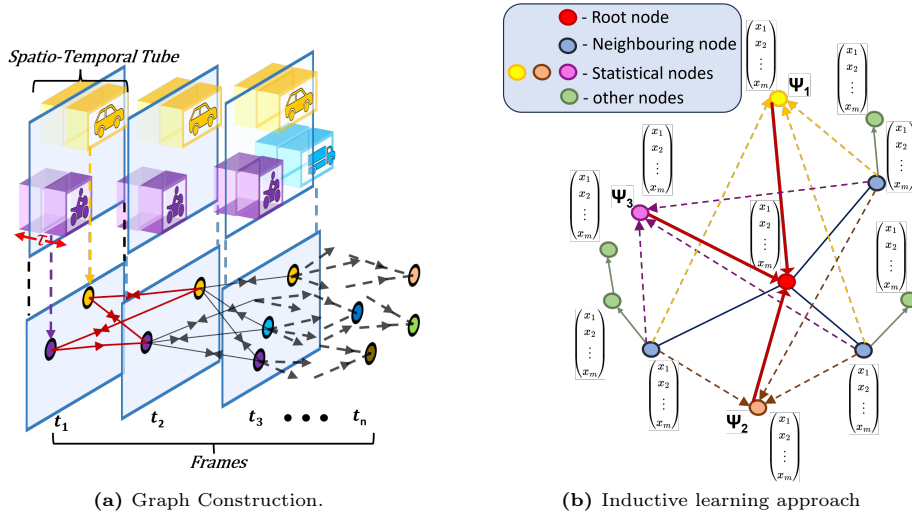


Fig. 1: (a) Temporal bidirectional bipartite graph construction from spatio-temporal tube features. (b) Inductive learning framework updating root node features via generated nodes through statistical aggregation from neighbors.

influenced by factors such as traffic signals, road conditions, and the collective behaviors of road users, making it challenging to accurately predict and recognize activities [22]. The diversity in driving behaviors and techniques across different regions and countries further complicates this task, highlighting the need for efficient algorithms capable of accurately interpreting a wide range of driving actions [31]. Addressing these challenges is vital for improving traffic safety and enhancing the effectiveness of autonomous driving technologies.

Traditional video activity recognition methods primarily focus on capturing the spatio-temporal dynamics of objects by tracking features like bounding boxes, 3D coordinates, and key points across video frames to understand how an object’s spatial and temporal attributes evolve over time [2, 16]. While these methods effectively capture broader spatio-temporal aspects, they often struggle with finer details, especially in cluttered environments where subtle movements are crucial [4]. To overcome these limitations, deep learning techniques such as Convolutional Neural Networks (CNNs) [1] and Long Short-Term Memory Networks (LSTMs) [25], along with more advanced approaches like 3D Convolutional Neural Networks (3D CNNs) [33], have been adopted. Unlike 2D CNNs that capture only spatial features from individual frames, 3D CNNs process both spatial and temporal data simultaneously, making them more effective at recognizing complex actions in dynamic scenes [33]. However, these models still face challenges, particularly with videos that require an understanding of extended temporal contexts spanning large video durations [10]. These methods typically process segmented video clips rather than entire sequences, limiting their ability to capture long-term dependencies. They also require substantial computational resources and may struggle to fully capture intricate relationships in dynamic

environments, crucial for understanding traffic behavior [28]. Our approach uses 3D CNNs to extract detailed spatio-temporal features from short object movements within continuous video streams. These features are connected through a our proposed graph structure, enabling efficient long-term video representation while managing computational demands, improving the model’s ability to represent complex traffic dynamics.

Graph-based approaches [24, 40] have become increasingly popular for tackling the complexities of traffic scene analysis by effectively modeling spatial and temporal relationships between objects. Spatio-temporal traffic graphs, for instance, capture the evolution of traffic states, enabling the identification of congestion patterns and flow dynamics in mixed traffic conditions [26]. Inductive learning techniques in graph-based models have further enhanced these approaches by enabling generalization to unseen data through feature-based embeddings, rather than relying solely on fixed graph structures [13, 37]. This is particularly crucial for real-time traffic analysis, where adaptability to dynamic environments is essential. Additionally, advanced aggregation techniques in Graph Neural Networks (GNNs) boost the representational power of models by effectively integrating information from neighboring nodes while maintaining computational efficiency [9]. By leveraging these advancements, our approach offers a scalable and robust solution for accurate activity recognition in complex driving scenarios.

In this work, we present a novel framework that utilizes inductive learning into graph-based approaches for driving scene analysis, inspired by the GraphSAGE [13] framework. We enhance the capabilities of GNNs by introducing a more effective node generation mechanism as part of the inductive learning process, incorporating a multi-aggregation approach [9]. This inductive approach allows our model to adapt to new scenarios without the need for retraining on a fixed graph structure, making it particularly suitable for real-time applications in dynamic and evolving traffic environments. We employ state-of-the-art (SotA) feature extraction techniques to capture robust spatio-temporal representations of objects within a very short time span. By incorporating high-order statistics into the aggregation process, we construct a high-order evolving graph that efficiently captures the intricate details of scene dynamics in video data. This ensures that our model remains accurate and robust, effectively handling the complexity of driving scenarios.

1.1 Related Work

Spatio-Temporal Approaches. Spatio-temporal methods in video analysis have been instrumental in advancing the understanding of object movements and interactions over time, which is crucial for tackling complex video tasks. Foundational models like Separable 3D CNNs (S3D) [36], R(2+1)D [34], Slow-Fast networks [11], and 3D ResNet [14] have been key in extracting rich spatio-temporal features. These architectures have improved motion representation and classification accuracy, setting benchmarks in processing efficiency and depth of analysis across various video applications, including action recognition and scene

understanding in dynamic environments. Building on these foundations, recent approaches have introduced more sophisticated mechanisms for refining video analysis. For instance, Diba *et al.* [8] proposed a convolution-attention network that balances short and long-range temporal cues using a hierarchical structure, providing nuanced temporal understanding necessary for complex video reasoning. Similarly, Chen *et al.* [6] utilized Spatio-Temporal Cross-Covariance Transformers in neural video compression, effectively merging spatial and temporal data with advanced 3D convolutional strategies and attention mechanisms, which are crucial for managing large-scale video data in autonomous driving applications. Li *et al.* [20] introduced the Spatio-Temporal Rationalizer (STR), a technique that dynamically identifies crucial video moments and objects, enhancing model performance on multifaceted datasets—an essential capability for accurate, context-aware traffic analysis in autonomous driving. Our work builds on these models as backbones, seamlessly integrating them into our GNN framework to significantly enhance scene classification, action recognition, and scene understanding in autonomous driving.

Graph-Based Approaches. Recent advancements in GNNs [21, 23, 30] for traffic scene analysis have shown substantial progress in traffic management and vehicle behavior prediction. Foundational work includes Visual Traffic Knowledge Graphs by Guo *et al.* [12], which parse complex interactions within traffic scenes to improve semantic understanding. Mlodzian *et al.* [23] contributed the nuScenes Knowledge Graph, which provides a comprehensive semantic representation that enhances trajectory prediction by incorporating detailed contextual information like road topology and traffic regulations. Further advancements include Thakur *et al.* [30], who developed a nested graph-based framework for early accident anticipation, contributing to proactive traffic safety. Li *et al.* [19] introduced EvolveGraph, a framework for multi-agent trajectory prediction using dynamic relational reasoning, which significantly improves the understanding and forecasting of agent movements in complex environments. Kumar *et*

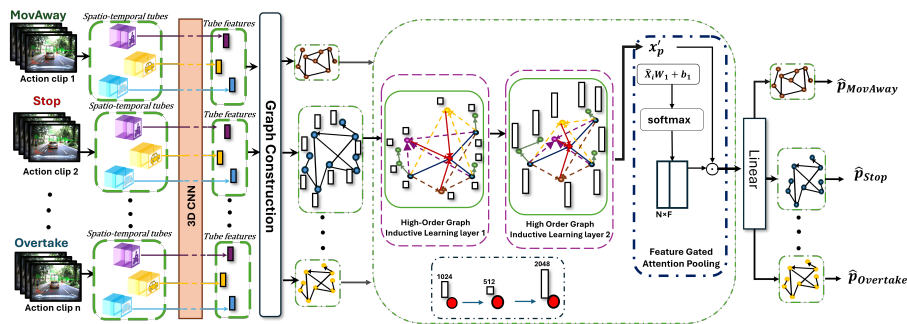


Fig. 2: The proposed pipeline begins with (a) graph construction, where nodes receive features extracted from spatio-temporal tubes via a 3D CNN. The graph is then processed through (b) high-order multi-aggregation based inductive learning layers, followed by (c) feature-level gated attention aggregation and a fully connected linear layer for classification.

al. [18] refined trajectory prediction with an interaction-based model that incorporates agent interactions over a hybrid traffic graph, enhancing accuracy in mixed traffic scenarios. Jin *et al.* [17] focused on fast contextual scene graph generation with unbiased context augmentation for dynamic scene understanding, while Hogle *et al.* [15] explored dynamic interaction-aware scene understanding crucial for reinforcement learning in autonomous driving. Additionally, Tian *et al.* [32] developed RSG-Search Plus, a traffic scene retrieval method based on road scene graphs, facilitating efficient scene indexing and retrieval. Liu *et al.* [21] combined traffic scenario understanding with video captioning, using a guidance attention captioning network to generate informative video summaries. Wang *et al.* [35] proposed a graph and spatio-temporal continuity framework for accident anticipation, highlighting GNNs’ utility in predicting traffic incidents. Sadid *et al.* [27] introduced a dynamic spatio-temporal GNN for surrounding-aware trajectory prediction in autonomous vehicles, and Zhou *et al.* [39] presented a hierarchical knowledge-guided traffic scene graph representation framework that supports intelligent vehicle applications by learning complex traffic scene dynamics. Büchner *et al.* [3] investigated the aggregation of lane graphs for urban automated driving, enhancing navigational strategies. Unlike these approaches, our method employs high-order evolving graphs based on inductive learning to better capture long-term dependencies and complex interactions, significantly boosting accuracy and robustness in scene understanding and action recognition for autonomous driving.

2 Proposed Approach

Our model pipeline, depicted in Fig. 2, classifies events in traffic scenes using a dataset of N videos, denoted as V_n where $n = \{1, 2, \dots, N\}$, each with a ground truth label p_n . The objective is to predict class probabilities \hat{p}_n that closely match the true labels p_n . The approach comprises three main components: 1) Spatio-temporal tube feature extraction using a 3D CNN backbone, 2) Construction of a temporal bidirectional bipartite graph linking objects across consecutive frames based on extracted features as node features, and 3) A multi-aggregation mechanism incorporates high-order statistics from neighboring nodes. This mechanism utilizes an inductive representation-based graph learning framework [13]. The process is finalized with feature-level gated pooling and a classification layer, producing the class probabilities \hat{p}_n to match the true labels p_n accurately.

2.1 Spatio-Temporal Tube Features

We start by generating spatio-temporal tubes for each object in every video frame. For each object a spatio-temporal tube is created based on its bounding box, capturing the object’s appearance across a sequence of frames. This tube spans τ consecutive frames, with $\tau/2$ frames preceding and $\tau/2$ frames following the central frame. The resulting spatio-temporal tube forms a 4D tensor with dimensions of τ (time) $\times w$ (width) $\times h$ (height) $\times C$ (channels), where $w \times h$ represents the bounding box size of the object in the central frame. We then adjust the spatial resolution of these tensors using bilinear interpolation to match

the input dimensions required by the pre-trained 3D CNN model. The S3D model [36], a SotA 3D CNN designed for video classification tasks, is employed as our feature extractor to process these spatio-temporal tubes. It is important to note that this feature extraction serves as a pre-processing step and does not involve any training. The extracted features are used as node features in our bipartite graph, modeling dynamic interactions between consecutive frames.

2.2 Temporal Bidirectional Bipartite Graph

Graph Construction. Given a video V_n consisting of T frames, we construct a multi-frame bidirectional bipartite graph $\mathcal{G}_n(\mathcal{V}_n, \mathcal{E}_n)$ to represent the interactions among objects between consecutive frames. The node set \mathcal{V}_n comprises all the objects present across entire video frames, while the edge set \mathcal{E}_n represents the bidirectional connections between objects in adjacent frames. Specifically, the graph \mathcal{G}_n connects objects represented by sets X_1, X_2, \dots, X_k , which are present in frames f_1, f_2, \dots, f_T of video V_n , through bidirectional bipartite connections. These connections are represented by the edge set $\mathcal{E}_n = B_{1 \leftrightarrow 2} \cup B_{2 \leftrightarrow 3} \cup \dots \cup B_{(T-1) \leftrightarrow T}$. Formally, this can be expressed as:

$$\mathcal{G}_n(\mathcal{V}_n, \mathcal{E}_n) = G \left(\bigcup_{i=1}^T X_i, \bigcup_{i=1}^{T-1} B_{i \leftrightarrow (i+1)} \right) \quad (1)$$

where $B_{i \leftrightarrow (i+1)}$ represents the bidirectional bipartite edge set between objects in frames i and $i + 1$. The total number of nodes and edges in the graph are therefore calculated as $\sum_{i=1}^T |X_i|$ and $2 \sum_{i=1}^{T-1} |X_i| \times |X_{i+1}|$, respectively. We uniformly sample frames at regular intervals, specifically every 5th frame (i.e., every 0.167 seconds for a 30fps video), to avoid constructing unnecessarily large graphs.

Complexity Analysis. The complexity of constructing our T -layered bipartite graph \mathcal{G}_n can be analyzed by evaluating the complexities associated with node and edge creation. For a video V_n with an average of \bar{x} objects per frame across T frames, the complexity of creating the nodes is approximately $O(T\bar{x})$. The creation of edges \mathcal{E}_n , which involves connecting objects between adjacent frames, results in a complexity of $O(T\bar{x}^2)$. Therefore, the overall complexity of constructing the graph, combining both node and edge creation, is $O(T\bar{x}) + O(T\bar{x}^2) \approx O(T\bar{x}^2)$, assuming $\bar{x} \gg 1$. In contrast, constructing a fully connected graph with $T \times \bar{x}$ nodes would lead to a complexity of approximately $O((T\bar{x})^2)$ for large T and \bar{x} , primarily because connections extend beyond just consecutive frames. This comparison highlights the efficiency of our bipartite graph approach in handling longer videos while keeping the average number of objects per frame constant. Notably, when $\bar{x} \approx 1$, the complexity of \mathcal{G}_n simplifies to $O(T)$, indicating that the graph’s complexity is more dependent on the number of frames than on the number of objects per frame.

2.3 High-Order Inductive Learning

The core concept of our approach is to efficiently learn and aggregate features (illustrated in Fig. 3) by generating robust node embeddings from a node’s local

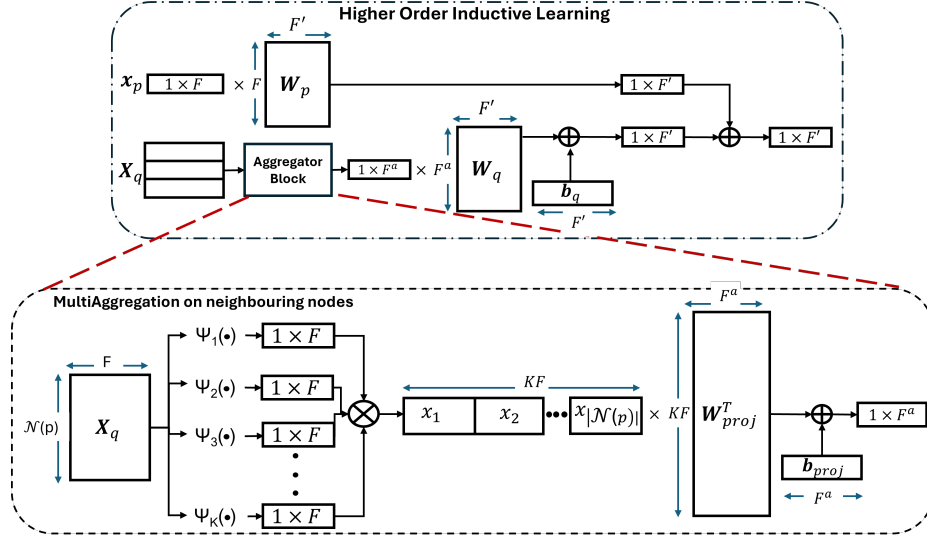


Fig. 3: Illustration of the multi-aggregation process in the high-order inductive learning mechanism as discussed in Section 2.3. The top block shows the graph convolution where the root node is updated through transformations and aggregation. The bottom block details the statistical multi-aggregation applied to neighboring nodes, with features concatenated and projected to enhance node representations.

neighborhood that generalize well to unseen interaction patterns in traffic scenes. Inspired by GraphSAGE [13], we extend its single aggregation operation by employing a set of K trainable aggregation functions, ψ_k for each $k \in \{1, \dots, K\}$, to aggregate information from neighbors. The aggregated features are then projected into a more informative embedding space using projection matrices \mathbf{W}_k , transforming the raw features into a richer, more expressive representation.

Generate Embedding and Update. The feature information of the p^{th} root node (denoted by \mathbf{x}_p) is updated based on its neighboring nodes indexed by $\mathcal{N}(p)$. This process begins by applying a set of K statistical aggregation functions ψ_k , each independently reducing the $|\mathcal{N}(p)| \times F$ feature matrix, constructed from the neighboring nodes, into a $1 \times F$ vector. These aggregated features are then concatenated along the feature dimension, resulting in a final dimension of $1 \times KF$. Without loss of generality, all the projection matrices $\mathbf{W}_k \in \mathbb{R}^{F \times F^a}$ associated with each aggregation function can be represented as a single projection matrix $\mathbf{W}_{proj} = [\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_K]^T \in \mathbb{R}^{KF \times F^a}$, where F^a is the output feature dimension of an aggregation operation. Additionally, a bias term $\mathbf{b}_{proj} \in \mathbb{R}^{1 \times F^a}$ is added to this transformation, which can be expressed as:

$$\Psi_{q \in \mathcal{N}(p)}(\mathbf{x}_q) = \mathbf{W}_{proj} \cdot \left(\left\| \psi_k(\mathbf{x}_q) \right\| \right) + \mathbf{b}_{proj}, \quad (2)$$

where q indexes the neighboring nodes relative to the root node. This final transformation updates the feature information of each node \mathbf{x}_p in our graph \mathcal{G}_n . The update rule for the l^{th} layer of our GNN model is therefore expressed

as:

$$\mathbf{x}_p^{l+1} = \mathbf{W}_p^l \mathbf{x}_p^l + \mathbf{W}_q^l \cdot \Psi_{q \in \mathcal{N}(p)}^l \mathbf{x}_q^l + \mathbf{b}_q^l, \quad (3)$$

Here, $\mathbf{W}_p \in \mathbb{R}^{F \times F'}$ and $\mathbf{W}_q \in \mathbb{R}^{F^a \times F'}$ are separate transformations applied to the root node and the neighboring nodes, respectively, to match the desired output dimension F' of the GNN operation.

Multi-Aggregation with High-Order Statistics. We enhanced the aggregation process by introducing high-order statistical moments alongside the usual mean, standard deviation, and median calculations, resulting in a more robust and informative aggregation operation. Each aggregator performs an element-wise operation along each feature dimension, reducing the $|\mathcal{N}(p)| \times F$ feature matrix \mathbf{X} into $1 \times F$ dimensional vector, where $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{|\mathcal{N}(p)}|]^T$ and each \mathbf{x}_i is an F -dimensional column vector. The mean aggregation function $\psi_1(\mathbf{X}) = \mu = \frac{1}{|\mathcal{N}(p)|} \sum_{i=1}^{|\mathcal{N}(p)|} \mathbf{X}_{i,:}$ computes the mean of feature \mathbf{X} . Similarly, the standard deviation of \mathbf{X} can be calculated as $\psi_2(\mathbf{X}) = \sigma^2 = \frac{1}{|\mathcal{N}(p)|} \sum_{i=1}^{|\mathcal{N}(p)|} (\mathbf{X}_{i,:} - \mu)^2$. The median aggregation function ψ_3 calculates the feature-wise median. For each feature position along $|\mathcal{N}(p)|$, the median is determined by first sorting the values in ascending order and then selecting the middle value, which corresponds to the position $i = \lfloor \frac{|\mathcal{N}(p)|}{2} \rfloor$. To extend the range of statistical properties captured by our model, we incorporate high-order statistical aggregations. These are represented as $\psi_m(\mathbf{X}) = \frac{1}{|\mathcal{N}(p)|} \sum_{i=1}^{|\mathcal{N}(p)|} (\mathbf{X}_{i,:} - \mu)^m$, where m denotes the order of the statistic. Specifically, we use high-order aggregations with $m = 3$ and $m = 4$, in addition to the mean, standard deviation, and median. This multi-aggregation strategy ensures that our model captures a diverse range of statistical properties from the node’s local neighborhood, leading to a more comprehensive and expressive node embedding.

Feature-Gated Attention Pooling. After generating the output features from the final layer of our GNN model, we apply feature-gated attention pooling to aggregate node features across the graph \mathcal{V}_n and produce the final prediction for the video sequence. This mechanism selectively emphasizes the most relevant nodes, enhancing the model’s ability to capture critical aspects of the video. Let the output of the final GNN layer be $\hat{\mathbf{X}} \in \mathbb{R}^{|\mathcal{V}_n| \times \hat{F}}$, where $|\mathcal{V}_n|$ is the number of nodes in the graph corresponding to the video V_n . The feature-gated attention pooling is expressed as:

$$\mathbf{X}_{pooled} = \sum_{i=1}^{|\mathcal{V}_n|} \mathcal{S}(\hat{\mathbf{X}}_i \mathbf{W}_1 + \mathbf{b}_1) \odot \hat{\mathbf{X}}_i \quad (4)$$

where term $\mathcal{S}(\cdot) = e^{(\cdot)} / \sum^{|\mathcal{V}_n|} e^{(\cdot)}$ is a softmax operation applied over $\hat{\mathbf{X}}_i$ using a learned transformation with weight matrices $\mathbf{W}_1 \in \mathbb{R}^{\hat{F} \times \hat{F}}$ and bias $\mathbf{b}_1 \in \mathbb{R}^{1 \times \hat{F}}$. This serves as a feature-gated attention mechanism, modulating the importance of each node’s features through an element-wise Hadamard product (denoted by \odot) with the input matrix $\hat{\mathbf{X}}$. By gating features, the pooling operation prioritizes the most relevant nodes, ensuring the model focuses on those contributing most

to the action representation. The pooled output \mathbf{X}_{pooled} is then passed through a fully connected layer followed by a softmax activation function to yield the final prediction \hat{p}_n for the video sequence V_n .

3 Experiments

In this section, we explore the datasets utilized, the evaluation metrics applied, and the implementation specifics. Following this, we present a performance comparison with existing SotA methods and conclude with an in-depth ablation analysis.

3.1 Dataset Description and Evaluation

Our method was evaluated using two primary datasets: the ROAD event Awareness Dataset (ROAD) Dataset [29] and the ROAD Waymo Dataset [10]. The ROAD is specifically tailored for autonomous driving research, comprising 122,000 frames extracted from 22 videos, each approximately 8 minutes long. The dataset was collected using a camera mounted on an autonomous vehicle (AV). It features two distinct sets of action labels. The first set, known as agent-level actions, describes the behaviors of various road agents (e.g., pedestrians, vehicles, cyclists) in a scene, such as “waiting2Cross” or “overtaking”, particularly in static contexts where the AV is not in motion. The second set, referred to as AV-actions, characterizes the autonomous vehicle’s own movements, independent of the actions of other agents in the scene, with labels such as “AV-on-the-move” and “AV-stopped”. The ROAD Waymo Dataset builds upon the ROAD dataset by incorporating data from the Waymo Open Dataset [10]. It consists of 198,000 annotated frames from 1,000 videos, each averaging 20 seconds in duration. This dataset offers action annotations for various road agents, along with agent type and semantic location labels. It contains 54,000 tracks, with over 3.9 million bounding box-level agent labels and more than 4.2 million action and location labels. The ROAD Waymo Dataset is designed for the robust evaluation of autonomous driving models, particularly in complex and diverse driving scenarios captured by Waymo’s autonomous vehicles. For evaluation, both datasets provide predefined splits for training, validation, and testing. Model performance was assessed using classification accuracy (Acc. in %) and mean Average Precision (mAP in %).

3.2 Implementation Details

Our implementation is built using PyTorch 2.12 and PyTorch Geometric 2.1. The proposed model architecture consists of two GNN layers that leverage multi-aggregation of high-order statistics from neighboring nodes to effectively process graph data. For each input video, we first extract spatio-temporal tube regions based on the bounding boxes of objects in the video frames. These regions are cropped across $\tau = 16$ frames, centered at a middle frame as described in Section 2.1. To ensure that the bounding boxes encompass the entire object appearance over the 16 frames, the bounding boxes are scaled by a factor of 1.5

before cropping. Each object instance is then resized to a fixed dimension of 256×256 , constructing a spatio-temporal tube represented as a 4D tensor of size $16 \times 256 \times 256 \times 3$. These tensors are fed into a Kinetics pre-trained 3D CNN feature extractor, specifically the S3D model [36], which outputs a 1024-dimensional vector for each spatio-temporal tube. These vectors serve as nodes in our constructed graph, with each node representing a distinct object. The node features are then processed through our proposed high-order statistics-based multi-aggregation GNN. To achieve a more meaningful and informative representation, we employ a bottleneck mechanism between the two GNN layers. In the first GNN layer, the input node feature dimension is reduced by half to 512. The output from this layer is then passed through a second GNN layer, where the feature dimension is restored to 1024. Next, we apply feature-gated attention pooling, as described in Section 2.3, directly to the 1024-dimensional output features. Since we use element-wise multiplication in the gating mechanism, the output dimension of the transformation weights is kept consistent at 1024, matching the output feature dimension. Our GNN model is trained using the Adam optimizer with a learning rate of 1×10^{-4} , a weight decay of 0.5, a batch size of 32 sequences, and for 200 epochs. During ablation studies with different backbone networks, which may have varying input node feature dimensions, we maintain the bottleneck feature dimension between the two GNN layers at a ratio of 0.5.

Table 1: Comparison of the results using features from different feature extractors on the agent tube features. The best results are shown in **bold**.

Agent Tube Features	S3D [36]	R(2+1)D [34]		SlowFast [11]		3D RES. [14]	
	Acc. mAP	Acc.	mAP	Acc.	mAP	Acc.	mAP
ROAD Dataset	55.87 34.43	60.42	34.80	56.44	30.26	47.69	19.26
ROAD Waymo Dataset	57.39 21.30	59.90	23.21	58.05	19.86	52.34	14.74

3.3 Performance Comparison

In this work, we establish a baseline by approaching the problem as a classification task on the ROAD and ROAD Waymo datasets, utilizing various video processing backbones. Table 1 compares the performance of four feature extraction methods like S3D [36], R(2+1)D [34], SlowFast [11], and 3D ResNet [14] on the ROAD and ROAD Waymo datasets, highlighting the effectiveness of our high-order GNN approach across different off-the-shelf backbones. R(2+1)D demonstrates the highest accuracy on both datasets, with 60.42% on ROAD and 59.90% on ROAD Waymo, showcasing its strong ability to correctly identify actions in driving scenes. However, its mAP scores, which measure the precision of ranking detections, decline significantly on the ROAD Waymo dataset (34.80% on ROAD and 23.21% on ROAD Waymo). This indicates that while R(2+1)D is effective at recognizing actions, it struggles with accurately ranking them in more complex scenarios. S3D provides a more balanced performance, achieving accuracies of 55.87% on ROAD and 57.39% on ROAD Waymo, with mAPs of 34.43% and 21.30%, respectively. Although S3D’s accuracy is slightly lower than

R(2+1)D’s, its more consistent mAP scores suggest it offers better precision across various actions, making it a reliable choice when both accuracy and precise ranking are important. SlowFast performs moderately well, with accuracies of 56.44% and 58.05%, but its mAP scores (30.26% and 19.86%) suggest it is less effective at accurately ranking detected actions. This may be due to its dual-pathway design, which might not fully align with the specific dynamics of these datasets. 3D ResNet shows the lowest performance, with accuracies of 47.69% on ROAD and 52.34% on ROAD Waymo, and mAPs of 19.26% and 14.74%. These results highlight its challenges in both identifying and accurately ranking actions in complex driving scenes. In summary, R(2+1)D is most effective for high accuracy in action identification, while S3D offers a balanced approach with reliable precision in action ranking. SlowFast and 3D ResNet may require further optimization to better handle the complexities of these driving scenarios.

Table 2 provides a comparison of action classification performance using Au-

Table 2: Performance comparison of S3D and R(2+1)D features computed for AV-actions for ROAD and ROAD Waymo datasets.

AV Features	S3D [36]		R(2+1)D [34]	
	Acc.	mAP	Acc.	mAP
ROAD	87.91	41.71	87.93	40.14
ROAD Waymo	93.71	36.16	93.14	30.42

tonomous Vehicle (AV) features, focusing on the interactions between multiple objects and the overall scene dynamics, across the ROAD and ROAD Waymo datasets. On the ROAD dataset, both S3D and R(2+1)D models perform exceptionally well, achieving nearly identical accuracies of 87.91% and 87.93%, respectively. However, S3D slightly outperforms R(2+1)D in terms of mAP, with scores of 41.71% versus 40.14%. This suggests that while both models are adept at recognizing driving actions, S3D is marginally better at ranking the detected actions, which could be crucial in scenarios requiring precise action ordering. On the ROAD Waymo dataset, which presents more complex and varied scenes, S3D continues to lead with an accuracy of 93.71% and an mAP of 36.16%, compared to R(2+1)D’s 93.14% accuracy and 30.42% mAP. The decrease in mAP for both models on this dataset underscores the challenge of accurately ranking actions in more diverse environments. Nevertheless, S3D’s superior mAP suggests it remains more reliable for tasks where the accurate ranking of detected actions is critical. Overall, these results demonstrate that S3D consistently provides better action ranking, particularly in complex environments, making it a strong candidate for scenarios requiring precision in both identification and ranking. R(2+1)D remains highly accurate in action identification, which is beneficial in contexts where correct detection is paramount.

3.4 Ablation Studies

Impact of varying aggregations. Table 3 analyzes the impact of different aggregation strategies within our GNN on the performance of action recognition in

Table 3: Accuracies and mAP for agent-level and AV-level actions on ROAD and ROAD Waymo datasets, using different aggregation combinations with S3D [36] and R(2+1)D [34] backbones. The best results are shown in **bold**.

Dataset	Aggregation combination(s)					S3D [36]				R(2+1)D [34]			
	Mean	Med.	Std	m=3	m=4	Agent		AV		Agent		AV	
						Acc.	mAP	Acc.	mAP	Acc.	mAP	Acc.	mAP
ROAD	✓					54.60	27.54	85.73	38.86	59.04	33.22	85.67	38.40
	✓	✓				54.90	28.56	86.08	39.90	59.35	33.35	85.97	38.76
	✓	✓	✓			55.82	28.70	86.20	40.95	59.50	33.94	86.12	38.84
	✓	✓	✓	✓	✓	55.45	30.31	86.79	41.35	60.27	34.20	86.44	39.63
	✓	✓	✓	✓	✓	55.87	34.43	87.91	41.71	60.42	34.80	87.93	40.14
Road Waymo	✓					52.97	18.78	92.80	30.39	56.91	21.78	90.53	28.26
	✓	✓				53.23	19.56	93.04	31.26	56.34	21.41	92.51	29.94
	✓	✓	✓			54.17	19.62	93.05	35.58	55.31	21.96	92.76	30.34
	✓	✓	✓	✓	✓	54.53	20.56	93.48	35.65	56.28	22.71	92.80	30.39
	✓	✓	✓	✓	✓	57.39	21.30	93.71	36.16	59.90	23.21	93.14	30.42

the ROAD and ROAD Waymo datasets, covering both agent-level and AV-level actions. The results show that applying multiple aggregation techniques significantly improves performance for both S3D and R(2+1)D feature extractors. For agent-level actions, both S3D and R(2+1)D feature extractors show notable improvements when multiple aggregation techniques are applied. On the ROAD dataset, R(2+1)D achieves its highest accuracy of 60.42% and mAP of 34.80% with a combination of mean, median, standard deviation, and high-order moments (m=3 and m=4). This combination allows the model to effectively capture central tendencies, variability, and subtle patterns within the spatio-temporal features. S3D also benefits from this approach, achieving an accuracy of 55.87% and mAP of 34.43%, highlighting the importance of diverse statistical measures in enhancing node representations. A similar trend is observed for AV-level actions. The most effective aggregation strategy for S3D, incorporating mean, median, standard deviation, and high-order moments, results in an accuracy of 87.91% and mAP of 41.71% on the ROAD dataset. This approach helps the GNN capture intricate patterns and subtle variations in the data, crucial for distinguishing different actions in complex traffic scenarios. On the ROAD Waymo dataset, the benefits of these aggregation strategies are even more evident. S3D achieves its highest performance with an accuracy of 57.39% and mAP of 21.30% for agent-level actions, and 93.71% accuracy with 36.16% mAP for AV-level actions. R(2+1)D similarly benefits, reaching 59.90% accuracy and 23.21% mAP for agent-level actions, and 93.14% accuracy with 30.42% mAP for AV-level actions. The use of these diverse aggregation strategies within the GNN framework enhances the model’s ability to represent spatio-temporal features and manage the complexities of dynamic driving scenarios. By capturing a broad range of statistical properties such as central tendencies, variability, and high-order patterns—the model is better equipped to differentiate between various actions and interactions, leading to improved overall performance in both agent-level and AV-level contexts.

Impact of different pooling mechanisms. The results in Table 4 demonstrate the effectiveness of various pooling and aggregation methods in traffic

Table 4: Impact of different pooling mechanisms on agent-level and AV-level actions on ROAD datasets across varying architectures. The results are shown in **bold**.

Pooling mechanism	S3D [36]				R(2+1)D [34]			
	Agent		AV		Agent		AV	
	Acc.	mAP	Acc.	mAP	Acc.	mAP	Acc.	mAP
Global Mean	51.53	27.38	84.78	40.20	53.22	24.78	85.37	38.83
Global Sum	54.14	29.44	86.08	40.65	54.29	27.15	85.82	38.30
Global Max	49.23	26.38	84.08	39.30	56.74	30.36	85.67	37.85
Feature-Gated Attn. (Ours)	55.87	34.43	87.91	41.71	60.42	34.80	87.93	40.14

scene analysis, particularly on the ROAD dataset. In these scenarios, the choice of pooling mechanism is crucial for capturing relevant spatio-temporal patterns, which directly impacts the model’s ability to interpret dynamic interactions. Global Mean pooling averages features across nodes, tends to oversimplify the diverse interactions in traffic scenes. This simplification often results in the loss of important information, as reflected in the lower performance metrics (e.g., 51.53% accuracy and 27.38% mAP for Agent-level action using S3D). Global Add pooling operation is more informative by summing features, can still miss subtle yet important variations in the data, leading to only modest improvements (e.g., 54.29% accuracy and 27.15% mAP for Agent-level action using R(2+1)D). Global Max pooling operation is more effective at capturing the most prominent features, fails to consider the nuanced interactions essential in complex environments like traffic scenes. By focusing solely on maximum values, this method overlooks the broader context of traffic dynamics, resulting in subpar performance (e.g., 49.23% and 56.74% for Agent-level action using S3D and R(2+1)D respectively). In contrast, our Feature-Gated Attention (Attn.) mechanism shows superior performance by dynamically prioritizing the most relevant features within the graph. This method is particularly beneficial in traffic scene analysis, where multiple agents interact in complex ways that require careful consideration of context and relevance. By selectively attending to critical features, Feature-Gated Attention significantly enhances the model’s ability to distinguish subtle variations in traffic behaviors and interactions. This leads to much higher accuracy and mAP scores across both S3D and R(2+1)D architectures (e.g., 55.87% accuracy and 34.43% mAP for Agent-level action using S3D, 87.91% accuracy and 41.71% mAP for AV-level action using S3D). This ablation highlights the effectiveness of attention mechanisms in GNNs for traffic scene analysis, where accurately capturing the intricate dynamics of road users is crucial for robust action recognition.

Table 5: Impact of feature dimension bottlenecking (Compression) between two GNN layers on performance across different datasets and architectures.

GNN Intermediate Feat. Dim.	S3D [36]				R(2+1)D [34]											
	Road		Road Waymo		Road		Road Waymo									
	Agent	AV	Agent	AV	Tube	AV	Tube	AV								
	Acc.	mAP	Acc.	mAP	Acc.	mAP	Acc.	mAP								
W/o Compression	48.62	23.83	82.31	37.33	54.53	20.57	93.69	36.10	54.29	27.31	83.57	36.14	53.63	20.28	92.74	29.76
W Compression (Ours)	55.87	34.43	87.91	41.71	57.39	21.30	93.71	36.16	60.42	34.80	87.93	40.14	59.91	23.21	93.14	30.42

Impact of feature compression. The results in Table 5 highlight the significant performance improvements achieved by applying feature compression (bottlenecking) between GNN layers across different datasets and architectures. For example, in the ROAD dataset using the S3D architecture, feature compression boosted the accuracy for Agent actions from 48.62% to 55.87% and mAP from 23.83% to 34.43%. Similarly, in the ROAD Waymo dataset, the accuracy and mAP for Agent actions increased from 54.53% to 57.39% and from 20.57% to 21.30%, respectively. These improvements are consistent across both S3D and R(2+1)D architectures and various action types, including agent and AV level actions, demonstrating that feature compression consistently enhances model performance. Technically, this enhancement is due to the bottleneck enforcing more efficient and discriminative feature representations, similar to the compression phase in autoencoders. By reducing the feature dimensions passed between layers, the GNN focuses on capturing the most critical aspects of the input data, which is vital for accurately recognizing complex spatio-temporal dynamics in traffic scenes. Moreover, feature compression helps mitigate overfitting, particularly with large and noisy datasets, by preventing the model from learning irrelevant correlations. Overall, applying feature compression between GNN layers results in a more efficient and robust model, improving generalization and performance in challenging environments like traffic scenes.

4 Conclusion

In conclusion, this paper presents a comprehensive framework that leverages High-Order Evolving Graphs and advanced GNNs to enhance the modeling and representation of traffic dynamics in autonomous driving applications. Our approach integrates temporal bidirectional bipartite graphs with 3D CNN features, multi-aggregation strategies, and attentional mechanisms as part of inductive learning. This combination effectively captures complex spatio-temporal interactions within dynamic traffic scenes, significantly improving both the accuracy and detail in traffic scene analysis. The effectiveness of our method is demonstrated through extensive experiments on the ROAD and ROAD Waymo datasets, where it establishes a comprehensive baseline for further developments in traffic behavior analysis. The results highlight the importance of integrating detailed spatio-temporal features with advanced graph-based techniques to achieve accurate and scalable traffic behavior analysis.

Acknowledgement: This research was partially funded by the UKRI EPSRC project ATTRACT (EP/X028631/1): A Trustworthy Robotic Autonomous System for Casualty Triage.

References

1. Abbas, Q., Ibrahim, M.E., Jaffar, M.A.: Video scene analysis: an overview and challenges on deep learning algorithms. *Multimedia Tools and Applications* **77**(16), 20415–20453 (2018) [2](#)
2. Behera, A., Chapman, M., Cohn, A.G., Hogg, D.C.: Egocentric activity recognition using histograms of oriented pairwise relations. In: *VISAPP*. vol. 2, pp. 22–30 (2014) [2](#)

3. Büchner, M., Zürn, J., Todoran, I.G., Valada, A., Burgard, W.: Learning and aggregating lane graphs for urban automated driving. In: CVPR. pp. 13415–13424 (2023) [5](#)
4. Chappa, N.V., Nguyen, P., Nelson, A.H., Seo, H.S., Li, X., Dobbs, P.D., Luu, K.: Spartan: Self-supervised spatiotemporal transformers approach to group activity recognition. In: CVPR. pp. 5158–5168 (2023) [2](#)
5. Chen, L., Wu, P., Chitta, K., Jaeger, B., Geiger, A., Li, H.: End-to-end autonomous driving: Challenges and frontiers. *IEEE Trans. Pattern Anal. Mach. Intell.* (2024) [1](#)
6. Chen, Z., Relic, L., Azevedo, R., Zhang, Y., Gross, M., Xu, D., Zhou, L., Schroers, C.: Neural video compression with spatio-temporal cross-covariance transformers. In: ACM MM. pp. 8543–8551 (2023) [4](#)
7. Chougule, A., Chamola, V., Sam, A., Yu, F.R., Sikdar, B.: A comprehensive review on limitations of autonomous driving and its impact on accidents and collisions. *IEEE Open Journal of Vehicular Technology* (2023) [1](#)
8. Diba, A., Sharma, V., Arzani, M., Van Gool, L., et al.: Spatio-temporal convolution-attention video network. In: ICCV. pp. 859–869 (2023) [4](#)
9. Ding, Y., Zhang, Z., Zhao, X., Hong, D., Li, W., Cai, W., Zhan, Y.: Af2gcn: Graph convolution with adaptive filters and aggregator fusion for hyperspectral image classification. *Information Sciences* **602**, 201–219 (2022) [3](#)
10. Ettinger, S., Cheng, S., Caine, B., Liu, C., Zhao, H., Pradhan, S., Chai, Y., Sapp, B., Qi, C.R., Zhou, Y., et al.: Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In: ICCV. pp. 9710–9719 (2021) [2](#), [9](#)
11. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: ICCV. pp. 6202–6211 (2019) [3](#), [10](#)
12. Guo, Y., Yin, F., Li, X.h., Yan, X., Xue, T., Mei, S., Liu, C.L.: Visual traffic knowledge graph generation from scene images. In: CVPR. pp. 21604–21613 (2023) [4](#)
13. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. *Advances in neural information processing systems* **30** (2017) [3](#), [5](#), [7](#)
14. Hara, K., Kataoka, H., Satoh, Y.: Towards good practice for action recognition with spatiotemporal 3d convolutions. In: ICPR. pp. 2516–2521 (2018) [3](#), [10](#)
15. Hügle, M., Kalweit, G., Werling, M., Boedecker, J.: Dynamic interaction-aware scene understanding for reinforcement learning in autonomous driving. In: ICRA. pp. 4329–4335 (2020) [5](#)
16. Jiang, B., Wang, M., Gan, W., Wu, W., Yan, J.: Stm: Spatiotemporal and motion encoding for action recognition. In: CVPR. pp. 2000–2009 (2019) [2](#)
17. Jin, T., Guo, F., Meng, Q., Zhu, S., Xi, X., Wang, W., Mu, Z., Song, W.: Fast contextual scene graph generation with unbiased context augmentation. In: CVPR. pp. 6302–6311 (2023) [5](#)
18. Kumar, S., Gu, Y., Hoang, J., Haynes, G.C., Marchetti-Bowick, M.: Interaction-based trajectory prediction over a hybrid traffic graph. In: IROS. pp. 5530–5535 (2021) [5](#)
19. Li, J., Yang, F., Tomizuka, M., Choi, C.: Evolvegraph: Multi-agent trajectory prediction with dynamic relational reasoning. *NeurIPS* **33**, 19783–19794 (2020) [4](#)
20. Li, Y., Xiao, J., Feng, C., Wang, X., Chua, T.S.: Discovering spatio-temporal rationales for video question answering. In: ICCV. pp. 13869–13878 (2023) [4](#)
21. Liu, C., Zhang, X., Chang, F., Li, S., Hao, P., Lu, Y., Wang, Y.: Traffic scenario understanding and video captioning via guidance attention captioning network. *IEEE Transactions on Intelligent Transportation Systems* (2023) [4](#), [5](#)

22. Mishra, S., Rajendran, P.K., Vecchietti, L.F., Har, D.: Sensing accident-prone features in urban scenes for proactive driving and accident prevention. *IEEE Transactions on Intelligent Transportation Systems* **24**(9), 9401–9414 (2023) [2](#)
23. Mlodzian, L., Sun, Z., Berkemeyer, H., Monka, S., Wang, Z., Dietze, S., Halilaj, L., Luetttin, J.: nuscenes knowledge graph—a comprehensive semantic representation of traffic scenes for trajectory prediction. In: *ICCV*. pp. 42–52 (2023) [4](#)
24. Monninger, T., Schmidt, J., Rupprecht, J., Raba, D., Jordan, J., Frank, D., Staab, S., Dietmayer, K.: Scene: Reasoning about traffic scenes using heterogeneous graph neural networks. *IEEE Robotics and Automation Letters* **8**(3), 1531–1538 (2023) [3](#)
25. Narayanan, A., Dwivedi, I., Dariush, B.: Dynamic traffic scene classification with space-time coherence. In: *ICRA*. pp. 5629–5635 (2019) [2](#)
26. Roy, D., Kumar, K.N., Mohan, C.K.: Defining traffic states using spatio-temporal traffic graphs. In: *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. pp. 1–6 (2020) [3](#)
27. Sadid, H., Antoniou, C.: Dynamic spatio-temporal graph neural network for surrounding-aware trajectory prediction of autonomous vehicles. *IEEE Transactions on Intelligent Vehicles* (2024) [5](#)
28. Schiappa, M.C., Biyani, N., Kamtam, P., Vyas, S., Palangi, H., Vineet, V., Rawat, Y.S.: A large-scale robustness analysis of video action recognition models. In: *CVPR*. pp. 14698–14708 (2023) [3](#)
29. Singh, G., Akrigg, S., Di Maio, M., Fontana, V., Alitappeh, R.J., Khan, S., Saha, S., Jeddisaravi, K., Yousefi, F., Culley, J., et al.: Road: The road event awareness dataset for autonomous driving. *IEEE transactions on pattern analysis and machine intelligence* **45**(1), 1036–1054 (2022) [9](#)
30. Thakur, N., Gouripeddi, P., Li, B.: Graph (graph): A nested graph-based framework for early accident anticipation. In: *WACV*. pp. 7533–7541 (2024) [4](#)
31. Thomas, E., McCrudden, C., Wharton, Z., Behera, A.: Perception of autonomous vehicles by the modern society: a survey. *IET Intell. Transp. Syst.* **14**(10), 1228–1239 (2020) [2](#)
32. Tian, Y., Carballo, A., Li, R., Thompson, S., Takeda, K.: Rsg-search plus: An advanced traffic scene retrieval methods based on road scene graph. In: *2024 IEEE Intelligent Vehicles Symposium (IV)*. pp. 1171–1178 (2024) [5](#)
33. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 4489–4497 (2015) [2](#)
34. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: *CVPR*. pp. 6450–6459 (2018) [3](#), [10](#), [11](#), [12](#), [13](#)
35. Wang, T., Chen, K., Chen, G., Li, B., Li, Z., Liu, Z., Jiang, C.: Gsc: A graph and spatio-temporal continuity based framework for accident anticipation. *IEEE Transactions on Intelligent Vehicles* **9**(1), 2249–2261 (2023) [5](#)
36. Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K.: Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 305–321 (2018) [3](#), [6](#), [10](#), [11](#), [12](#), [13](#)
37. Yang, Z., Cohen, W., Salakhudinov, R.: Revisiting semi-supervised learning with graph embeddings. In: *International conference on machine learning*. pp. 40–48 (2016) [3](#)

38. Yurtsever, E., Lambert, J., Carballo, A., Takeda, K.: A survey of autonomous driving: Common practices and emerging technologies. *IEEE access* **8**, 58443–58469 (2020) [1](#)
39. Zhou, Y., Liu, X., Guo, Z., Cai, M., Gou, C.: Hktsq: A hierarchical knowledge-guided traffic scene graph representation learning framework for intelligent vehicles. *IEEE Transactions on Intelligent Vehicles* (2024) [5](#)
40. Zhou, Y., Zhang, Y., Zhao, Z., Zhang, K., Gou, C.: Toward driving scene understanding: A paradigm and benchmark dataset for ego-centric traffic scene graph representation. *IEEE Journal of Radio Frequency Identification* **6**, 962–967 (2022) [3](#)