

BP-SGCN: Behavioral Pseudo-Label Informed Sparse Graph Convolution Network for Pedestrian and Heterogeneous Trajectory Prediction

Ruochen Li¹, Stamos Katsigiannis¹, *Member, IEEE*, Tae-Kyun Kim¹, Hubert P. H. Shum^{1†}, *Senior Member, IEEE*

Abstract—Trajectory prediction allows better decision-making in applications of autonomous vehicles or surveillance by predicting the short-term future movement of traffic agents. It is classified into pedestrian or heterogeneous trajectory prediction. The former exploits the relatively consistent behavior of pedestrians, but is limited in real-world scenarios with heterogeneous traffic agents such as cyclists and vehicles. The latter typically relies on extra class label information to distinguish the heterogeneous agents, but such labels are costly to annotate and cannot be generalized to represent different behaviors within the same class of agents. In this work, we introduce the behavioral pseudo-labels that effectively capture the behavior distributions of pedestrians and heterogeneous agents solely based on their motion features, significantly improving the accuracy of trajectory prediction. To implement the framework, we propose the Behavioral Pseudo-Label Informed Sparse Graph Convolution Network (BP-SGCN) that learns pseudo-labels and informs a trajectory predictor. For optimization, we propose a cascaded training scheme, in which we first learn the pseudo-labels in an unsupervised manner, and then perform end-to-end fine-tuning on the labels in the direction of increasing the trajectory prediction accuracy. Experiments show that our pseudo-labels effectively model different behavior clusters and improve trajectory prediction. Our proposed BP-SGCN outperforms existing methods using both pedestrian (ETH/UCY, pedestrian-only SDD) and heterogeneous agent datasets (SDD, Argoverse 1).

Index Terms—Trajectory prediction, pedestrian, heterogeneous agents, behavioral pseudo-label, graph convolutional networks

I. INTRODUCTION

PREDICTING the future movement of traffic agents, known as trajectory prediction, is crucial for safe and efficient decision-making in applications such as autonomous vehicles [1]. Thanks to reliable data-driven [2] object tracking methods [3], accurate geometric trajectories can be extracted from videos, serving as a more representative feature set for modeling. Graph Convolutional Networks (GCNs) [4] have shown exceptional performance across diverse fields due to their adeptness at capturing spatial relationships [5]–[9]. This enables them to excel in applications ranging from trajectory agent interaction modeling [10]–[13] to human skeleton-based behavior modeling [14]–[18], highlighting the superior capabilities in handling graph-based data structures. Similarly, recognizing distinct movement behavior patterns among agents is pivotal to model




Fig. 1. We propose the behavioral pseudo-labels learned from observed trajectories, effectively representing inter- and intra-type behavioral differences to improve pedestrian and heterogeneous trajectory prediction accuracy.

the temporal dependency [19]. These patterns, when integrated with GCN, further enhance the precision of predictions by accounting for the inherent behavioral tendencies.

Existing trajectory prediction methods can be broadly classified into two categories. The first focuses on predicting *pedestrian trajectories* in datasets that are exclusively composed of pedestrians [20], [21] or deliberately omit non-pedestrian traffic agents [22]–[24]. These methods primarily employ neural networks to account for pedestrian social interactions, such as the pooling window mechanism [25] and social interaction graphs [10]–[12]. The second category encompasses *heterogeneous trajectory* prediction, considering a diverse range of traffic agents (e.g. cars, cyclists, pedestrians, etc.). Recent methods [26]–[28] exploit the annotated class labels of traffic agents to better model agent interactions in intricate urban scenarios. These labels facilitate the system’s understanding on multifaceted interactions among various agent types [28].

A notable research gap can be observed between pedestrian-only and heterogeneous trajectory prediction. Methods tailored solely for pedestrian behavior excel due to its predictable patterns but lack applicability in real-world scenarios like autonomous driving, since pedestrians behave very differently from heterogeneous agents [26], [28]. The fundamental differences in modeling the motion patterns of different types of agents stem from their distinct dynamics, speed ranges, spatial needs, interaction behaviors, decision-making processes, and ways of perceiving the environment, necessitating varied modeling approaches to accurately predict their trajectories. For heterogeneous trajectory prediction, ground-truth (GT) labels for agent types have traditionally been used to guide discriminative learning [26]–[29]. However, these labels often fail to capture diverse within-class behaviors: for example, ‘vans’ and ‘compact cars’ are both labeled simply as ‘cars,’ while ‘pedestrians’ can range from ordinary walkers to skate-

R. Li, S. Katsigiannis and H. P. H. Shum are with Durham University, UK. (e-mail: {ruochen.li, stamos.katsigiannis, hubert.shum}@durham.ac.uk).

T. K. Kim is with KAIST, Korea. (e-mail: kimtaekyun@kaist.ac.kr).

† Corresponding author: H. P. H. Shum

boarders [30]. This granularity issue can lead to mislabeling, especially when visually similar categories are grouped together. Moreover, obtaining such detailed GT labels is time-consuming and expensive. We argue that purely relying on manual labels is both insufficient and cost-ineffective for representing the nuanced motion patterns seen in real-world traffic scenarios.

In this paper, we present a unified framework utilizing machine-learned behavioral pseudo-labels applicable to both heterogeneous and exclusively pedestrian domains. Our insight is that behavioral pseudo-labels can capture both inter-class and intra-class behavioral variations among agents, thereby improving the accuracy of our model. For heterogeneous scenarios, the use of behavioral pseudo-labels eliminates the need for manual label annotations, streamlining the process and reducing the reliance on extensive labeled datasets. In pedestrian-only scenarios, these pseudo-labels facilitate the differentiation and learning of intrinsic motion patterns among pedestrians, offering a more nuanced understanding of pedestrian behavior. A shared advantage across both contexts is the significant improvement in overall prediction performance, demonstrating the versatility and efficacy of behavioral pseudo-labels in diverse trajectory prediction tasks (Fig. 1).

We propose the Behavioral Pseudo-Label informed Sparse Graph Convolution Network (BP-SGCN) for pedestrian and heterogeneous trajectory prediction. The network includes two modules. First, we introduce a deep unsupervised behavior clustering module that assigns pseudo-labels to agents based on their observed trajectories. This module marks a novel application of deep embedded clustering [31], utilizing high-level temporal latent features. It is supported by a Variational Recurrent Neural Network (VRNN) [32] that processes a set of customized geometric features, crucial for capturing motion dynamics such as speed, angle, and acceleration. Additionally, a soft dynamic time warping loss addresses temporal variances in trajectories, uniquely tailoring our approach for trajectory modeling. The generated behavioral pseudo-labels are specifically designed to enhance trajectory forecasting, highlighting our model's focus on the nuanced demands of trajectory prediction in complex environments. Second, we propose a goal-guided pseudo-label informed trajectory prediction module, which adapts SGCN [10], a powerful GCN backbone for trajectory prediction that utilizes a sparse spatial-temporal attention mechanism to effectively model spatial interactions and temporal dependencies of agents. We then employ a Gumbel-Softmax straight-through estimator to link up the clustering module, allowing the prediction module and clustering module to be fine-tuned in an end-to-end manner. Finally, we design a cascaded training scheme [33] that first trains pseudo-label clustering in an unsupervised manner, and then fine-tunes both clustering and trajectory prediction together with the prediction loss to maximize their compatibility.

BP-SGCN surpasses SOTAs in both heterogeneous prediction on the SDD [30] and Argoverse 1 [34] datasets, and in pedestrian prediction on the ETH/UCY [20], [21] dataset and the pedestrian-only setup of SDD [35]. Our source code is available at <https://github.com/Carrotsniper/BP-SGCN> to facilitate further research. Our contributions are:

- We propose the novel concept of behavioral pseudo-

labels to represent clusters of traffic agents with different movement behaviors, improving trajectory prediction without the need for any extra annotation.

- To implement the idea, we propose BP-SGCN, which introduces a cascaded training scheme to optimize the compatibility of its two core modules: the pseudo-label clustering module and the trajectory prediction module.
- We propose a deep unsupervised behavior clustering module to obtain behavioral pseudo-labels, tailoring the geometric feature representation and the loss to best learn the agents' behaviors.
- We propose a pseudo-label informed goal-guided trajectory prediction module, which facilitates end-to-end fine-tuning with its prediction loss for better clustering and prediction, outperforming existing pedestrian and heterogeneous prediction methods.

II. RELATED WORK

A. Trajectory Prediction

Deep learning models have driven the latest advancement of trajectory prediction. Social-LSTM [25] introduces RNN-based neural networks [36] to model the trajectories of pedestrians and a pooling window mechanism to describe the interactions among them. Social-GAN [37] incorporates the ideas of Generative Adversarial Networks [38] to predict multiple multi-modal trajectories with distance-based interaction modeling. Diffusion models [39] are adopted into trajectory prediction [40], [41], showing significant improvement.

Graph representations are increasingly recognized for their prowess in modeling relational features. TrafficPredict [42] incorporates soft-attention-based interaction graphs with LSTM to represent social interactions. STGAT [11] models the trajectories using Spatial-Temporal Graph Attention Networks based on the sequence-to-sequence architecture. Social-STGCNN [12] introduces weighted graph edges, providing an interpretable measurement of pedestrian interactions. STAR [43] takes advantage of the Transformer [44] to construct a spatial-temporal graph transformer for trajectory representation. SGCN [10] advances this by proposing sparse directed spatial-temporal graph representations to model spatial interactions and motion tendencies for each pedestrian. However, these methods primarily focus on modeling pedestrian interactions, overlooking the intricate interactions among heterogeneous agents.

Future trajectory or goal information enhances the trajectory prediction as it provides valuable insights into the long-term intentions of individual agents [19], [22]–[24], [40], [45]–[49]. CVAE [50] based methods [45]–[47] employ the future and past trajectory encoder during the training phase to train the latent representation for each agent. Such latent is used to generate future trajectories during inference. Goal retrieval methods sample goal points and incorporate them as guides to model the prediction diversity [22]–[24]. We employ the goal retrieval approach proposed in [23], as it does not require training a separate models as in the CVAE-based method.

As trajectories are influenced by their surroundings, some studies employ image contextual features as auxiliary information. Convolutional Neural Networks (CNN) are applied

to extract such features to improve prediction accuracy [51], [52]. Semantic segmentation is adapted on scene images to extract annotated image features [22]. While acknowledging the potential benefits of these techniques, we do not utilize scene-based features in this research. The integration of such features typically requires an auxiliary image processing network [52]–[54], which complicates the model architecture and detracts from our primary focus on exploring the impact of behavioral pseudo-labels. Our approach, although not incorporating scene-based data, is conceptually aligned with these methods in its attempt to capture the nuanced behaviors and interactions of traffic agents purely from trajectory data. The proposed method is orthogonal and complementary to the use of scene-based features. The two methods can be combined in future work.




Fig. 2. Trajectory visualization on heterogeneous SSD dataset, where red, green and blue dots represent pedestrians, bikers and cars, respectively. (a) and (c) represent heterogeneous scenarios with all agent types, (b) and (d) represent the pedestrian-only scenarios commonly used by pedestrian trajectory predictions [22], [24] by simply removing all non-pedestrian agents.

B. Heterogeneous Trajectory Prediction

Heterogeneous trajectory prediction considers traffic agents of all types. Adopting prediction approaches [13], [22], [46] by simply ignoring non-pedestrian agents, or considering all agents to be of the same class, results in sub-optimal performance. Heterogeneous methods focus on modeling different agent behaviors. VP-LSTM [55] separately treats vehicles and pedestrians with LSTMs. Proposal-based approaches such as CoverNet [56] generate predefined multimodal trajectory anchors from observations of both vehicles and pedestrians. For better interpretability and representation of heterogeneous agent interactions, graph-based attention mechanisms [11], [44] are proposed. NLNI [27] presents a novel spatial-temporal category graph and proposed graph attention to capture the category-wise and agent-wise interactions. Multiclass-SGCN [28] and Semantic-STGCNN [26] introduce one-hot encoding to encode annotated class labels as part of node features. HIMRAE [57] proposes dynamic interaction graphs among agents to reduce accumulated error for heterogeneous trajectory prediction. SMGCN [58] introduces a sparse multi-relational GCN to learn heterogeneous interactions among agents.

However, while these approaches achieve superior performance, most methods rely heavily on ground-truth labels to distinguish agent types. On the one hand, manual labeling is costly and error-prone, making the model’s performance overly dependent on label quality. On the other hand, focusing on semantic labels can overlook subtle behavioral differences among agents with the same semantic label. In this work, we propose learning behavioral pseudo-labels from agent motion dynamics, thereby reducing reliance on manual labels and capturing a broader spectrum of behaviors.

C. Motion Behavior Clustering

The clustering of temporal trajectory patterns allows modeling the behavioral groups for better trajectory prediction [59], [60]. Early works focus on the raw trajectory represented as 2D coordinates. Support vector clustering is introduced as a closed-loop method on motion vectors for motion behavior representations [61]. K-means on trajectory vectors or sequence key points obtain cluster centers to enhance trajectory prediction [60], [62]. DBSCAN is proposed to avoid manually specifying cluster numbers, adding more flexibility and interpretability to behavior patterns [59]. GP-Graph directly uses the absolute distance among pedestrians to determine the division of group [63]. The recent PCCSNet leverages BiLSTM network to encode coordinates prior to K-means clustering, identifying behavioral modalities [19]. In addition to modalities, FEND further applies 1D CNN and LSTM for trajectory encoding and employs the K-means for long-tail trajectory clustering to distinguish trajectory patterns [64].

However, most existing methods rely on shallow trajectory representations, limiting their ability to capture nuanced, evolving behaviors. Additionally, distance-based clustering approaches often struggle with complex motion patterns. To address these issues, we propose a cascaded optimization scheme featuring an end-to-end Deep Embedded Clustering (DEC) [31] module, which iteratively refines cluster assignments using a KL-divergence objective. This dynamic adaptation yields richer latent representations, enabling a more data-driven and expressive approach to modeling agent behaviors.

III. BEHAVIOR PSEUDO-LABEL INFORMED SPARSE GRAPH CONVOLUTION NETWORK

A. The High-Level Network Architecture

We observe a research gap in pedestrian and heterogeneous trajectory prediction. Existing pedestrian prediction approaches have limited applicability to heterogeneous traffic agents due to the diverse behaviors of agents. For instance, in Fig. 2, (a) and (c) depict intricate heterogeneous scenarios with bikers and cars exhibiting longer, non-linear paths, while pedestrian-only scenarios (b) and (d) overlook interactions among pedestrians, bikers and cars.

Although introducing annotated class labels for heterogeneous agents leads to better prediction performance [26]–[28], such labels are only a proxy of movement behaviors, which cannot represent intra-class behavioral differences and inter-class behavioral similarity.

To this end, we present the concept of *behavioral pseudo-labels*, which capture movement behaviors to enhance trajectory prediction. Our pseudo-labels do not require annotations, mitigating the risk of mislabeling and reducing labor costs. It can be applied to both pedestrian-only and heterogeneous datasets, resulting in superior prediction performance.

To realize pseudo-label informed trajectory prediction, we propose the *Behavioral Pseudo-Label Informed Sparse Graph Convolution Network* (BP-SGCN). As shown in Fig. 3, BP-SGCN includes two modules: deep unsupervised clustering and pseudo-label informed trajectory prediction. The former learns the pseudo-labels in an unsupervised manner, while the

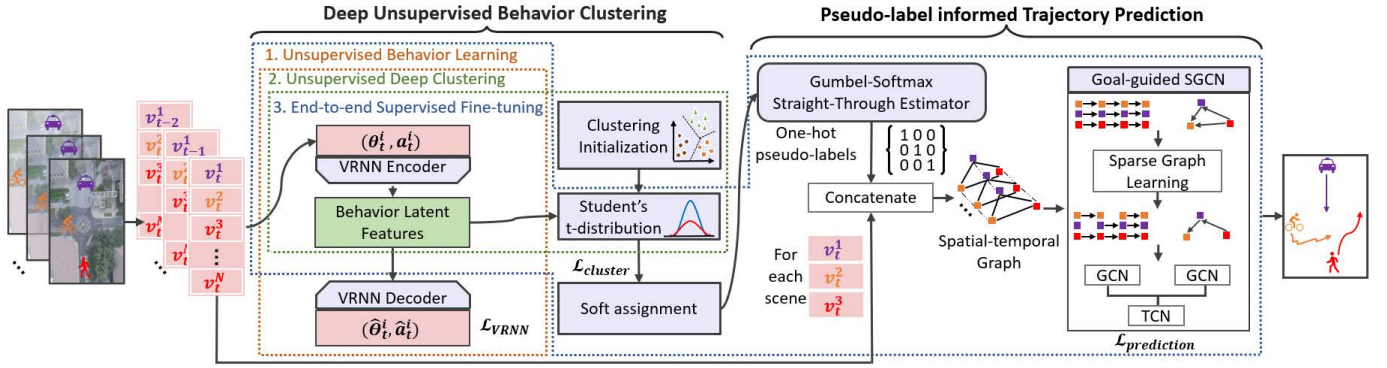


Fig. 3. The overview of BP-SGCN to learn the pseudo-labels for trajectory prediction, consisting of the deep unsupervised clustering module and the pseudo-label informed trajectory prediction module. We propose a cascaded optimization scheme to first learn pseudo-labels in an unsupervised manner, and then fine-tune them in an end-to-end manner with trajectory prediction supervision.

latter performs end-to-end optimization to improve pseudo-label clustering while predicting trajectories with such labels.

We propose the *cascaded training scheme* to obtain the pseudo-labels and thus high-quality trajectory prediction. First, highlighted with the orange dotted block in Fig. 3, the unsupervised behavior representation learning module derives behavior latent representations from observed trajectories through a Variational Recurrent Neural Network (VRNN) [32] module. Then, in the green dotted block, the behavior latent representations are fed into simple clustering modules (e.g., K-means, GMM, etc.) for cluster center initialization. We then perform unsupervised deep clustering to learn the distribution of pseudo-labels by feeding the VRNN latent representations to the Student’s t-distribution kernel [65]. This allows fine-tuning the VRNN encoder to create a better latent space and refine the cluster centers. Finally, indicated by the blue dotted block, we utilize a Gumbel-Softmax straight-through estimator to sample one-hot pseudo-labels, which are concatenated to the trajectory features as the input of goal-guided SGCN [10] for trajectory prediction. The whole network is optimized end-to-end, fine-tuning the pseudo-label clustering module to maximize its compatibility for trajectory prediction.

B. Deep Unsupervised Behavior Clustering

Here, we explain how we obtain behavior clusters, which serve as powerful features for effective trajectory prediction.

1) Geometric Representation of Trajectories

Given a series of observed video frames of N agents over time $t \in [1, T_{obs}]$, and the corresponding 2-D trajectory coordinates (x_t^i, y_t^i) , $i \in [1, N]$, our objective is to predict the future trajectory coordinates $p_t^i = (x_t^i, y_t^i)$ of each traffic agent i within a time horizon $t \in [T_{obs}+1, T_{pred}]$.

We introduce relative angle and acceleration magnitude to learn behavior latents. While global velocity is an effective feature for trajectory prediction [10], [28], it is less representative of behaviors, as it depends on global movement directions, and is less sensitive to velocity changes. Relative angles provide a representation that is invariant to the initial facing direction,

which is complemented with the magnitude of acceleration that has been shown to be effective for modeling behaviors.

For each traffic agent i , we calculate its velocity vector at time t . For simplicity, we remove the notation i in the following equation:

$$\mathbf{v}_t = \left(\frac{x_t - x_{t-1}}{t - (t-1)}, \frac{y_t - y_{t-1}}{t - (t-1)} \right), \quad (1)$$

where $\forall t \in [1, T_{obs}]$, we compute the cosine of the angle, $\cos(\theta_t)$ between velocity vectors, \mathbf{v}_t and \mathbf{v}_{t-1} :

$$\cos(\theta_t) = \frac{\mathbf{v}_t \cdot \mathbf{v}_{t-1}}{|\mathbf{v}_t| \cdot |\mathbf{v}_{t-1}|}, \quad (2)$$

and the magnitude of corresponding acceleration at time t :

$$|\mathbf{a}_t| = \left| \frac{\mathbf{v}_t - \mathbf{v}_{t-1}}{t - (t-1)} \right| \quad (3)$$

The geometric feature is constructed as $g_t = (\cos(\theta_t), |\mathbf{a}_t|)$.

2) Behavior Representation Learning

We adapt VRNN to learn latent representations for behavior clustering [31], [66]. VRNN learns the temporal dependencies of a sequence by modeling the distribution over its hidden states with an encoder-decoder architecture. Compared to LSTM-based autoencoders [66], it effectively models the highly nonlinear dynamics and captures the uncertainties of latent space. Its probabilistic nature of variational inference improves the learning of implicit sequential data distributions.

In particular, the encoder network $\varphi_{enc}(\cdot, \cdot)$ receives the embedded geometric data $\varphi^g(g_t)$ and recurrent hidden state h_{t-1} to approximate the posterior distribution $q_\phi(\cdot)$:

$$\begin{aligned} q_\phi(z_t | g_{\leq t}, z_{< t}) &= \mathcal{N}(z_t | (\mu_{z,t}, \sigma_{z,t}^2)), \\ [\mu_{z,t}, \sigma_{z,t}] &= \varphi_{enc}(\varphi^g(g_t), h_{t-1}), \end{aligned} \quad (4)$$

where z_t is sampled using a reparameterization trick [67]. The decoder network $\varphi_{dec}(\cdot, \cdot)$ takes the embedded latent $\varphi^z(z_t)$ and h_{t-1} to approximate the reconstruction distribution $p_\delta(\cdot)$:

$$\begin{aligned} p_\delta(g_t | z_{\leq t}, g_{< t}) &= \mathcal{N}(g_t | (\mu_{g,t}, \sigma_{g,t}^2)), \\ [\mu_{g,t}, \sigma_{g,t}] &= \varphi_{dec}(\varphi^z(z_t), h_{t-1}). \end{aligned} \quad (5)$$

To enhance the temporal dependencies in sequences, the prior distribution in VRNN relies on h_{t-1} with $\varphi_{prior}(\cdot)$:

$$\begin{aligned} p_\delta(z_t|z_{<t}, g_{<t}) &= \mathcal{N}(z_t | (\mu_{0,t}, \sigma_{0,t}^2)), \\ [\mu_{0,t}, \sigma_{0,t}] &= \varphi_{prior}(h_{t-1}). \end{aligned} \quad (6)$$

We employ the Gated Recurrent Unit (GRU) [68] to update the RNN hidden state, which outperforms LSTM [36] when the sequence length is relatively short:

$$h_t = GRU(\varphi^g(g_t), \varphi^z(z_t), h_{t-1}). \quad (7)$$

The VRNN is optimized with a customized loss:

$$\mathcal{L}_{VRNN} = \mathcal{L}_{Soft-DTW} + \mathcal{L}_{ELBO}, \quad (8)$$

where $\mathcal{L}_{Soft-DTW}$ is a differentiable soft Dynamic Time Warping (DTW) loss [69]:

$$\mathcal{L}_{Soft-DTW} = \min_{\mu_{g,t}} \sum_{i=1}^N \frac{1}{T_{obs}} DTW_\gamma(\mu_{g,t}, g_t), \quad (9)$$

DTW_γ refers to the original DTW [70] discrepancy that measures and aligns the similarity between two time series, γ is a parameter indicating the acceptable distortion for aligning two sequences, $\mu_{g,t}$ is the decoded mean of the VRNN decoder. The loss allows capturing non-linear temporal alignment [71], which cannot be achieved with MSE. \mathcal{L}_{ELBO} is the variational evidence lower-bound with the Kullback–Leibler (KL) divergence [32], [67]:

$$\begin{aligned} \mathcal{L}_{ELBO} &= \mathbb{E}_{q_\phi(z_{\leq T_{obs}} | g_{\leq T_{obs}})} \left[\sum_{t=1}^{T_{obs}} (\log p_\delta(g_t | z_{\leq t}, g_{<t})) \right. \\ &\quad \left. - KL(q_\phi(z_t | g_{\leq t}, z_{<t}) || p_\delta(z_t | z_{<t}, g_{<t})) \right]. \end{aligned} \quad (10)$$

By optimizing \mathcal{L}_{VRNN} , the model aligns predicted and observed sequences while maintaining a theoretically grounded variational framework. This alignment enhances flexibility in handling non-linear temporal dynamics, and the KL regularization constrains the latent structure, thus ensuring stable training. Consequently, the VRNN encoder provides richer latent representations for subsequent unsupervised deep clustering, effectively leveraging spatio-temporal structures to capture nuanced agent behaviors.

3) Deep Embedded Clustering

We present a new application of Deep Embedded Clustering (DEC) [31] to cluster the agent behaviors latents from the VRNN encoder, thereby generating a distribution of pseudo-labels. DEC allows jointly optimizing the cluster centers and the VRNN encoder, enhancing the latent representation via back-propagation. This significantly outperforms traditional methods like k-means [72] and Gaussian mixture models [73], which lack the capability to refine input feature representations.

The initial phase of DEC involves setting cluster centers using VRNN behavior latents. We input all training data into the VRNN encoder to obtain the set of behavior latent features \mathbb{Z} , and then apply k-means to determine initial centers, $c_j \in [1, k]$.

Given the variance in agent behaviors across datasets, k is an empirically tuned hyperparameter.

We then apply Student’s T-Distribution [65], that is, Q distribution to compute the soft assignment between each initialized cluster center and latent vector [31]. Its kernel measures the probability of each encoded vector $z_i \in \mathbb{Z}$ belonging to the cluster j :

$$q_{ij} = \frac{\left(1 + \frac{d(z_i, c_j)}{\alpha}\right)^{-\frac{\alpha+1}{2}}}{\sum_{j'} \left(1 + \frac{d(z_i, c_{j'})}{\alpha}\right)^{-\frac{\alpha+1}{2}}}, \quad (11)$$

where d is a similarity metric that refers to the distance between the encoded vector z_i and center c_j , and α is the number of degrees of freedom of the Q distribution. We denote d as the Euclidean distance and set α to 1.

Meanwhile, we optimize the clustering network with a KL divergence loss to minimize the discrepancy between the two distributions:

$$\mathcal{L}_{cluster} = KL(P||Q) = \sum_i \sum_j \left(p_{ij} \log \frac{p_{ij}}{q_{ij}} \right), \quad (12)$$

where P is the auxiliary distribution:

$$p_{ij} = \frac{q_{ij}^2 / f_j}{\sum_{j'} q_{ij'}^2 / f_{j'}}, \quad (13)$$

and $f_j = \sum_i q_{ij}$ are soft cluster frequencies. Here, P is re-weighted from Q distribution in a way that sharpens high-confidence assignments and de-emphasizes low-confidence ones [31], thereby systematically increasing the separation between clusters in the latent space. Finally, we derive soft assignments from the Student’s t-distribution, reflecting the probability of the latent z_i^j in each cluster c_j . This approach not only offers greater flexibility in representing complex behaviors but also sharpens cluster boundaries by reinforcing high-confidence assignments and reducing ambiguity in low-confidence ones. Consequently, it yields more coherent clusters and better captures the inherent diversity in agent dynamics, ultimately enhancing the overall clustering quality.

C. Pseudo-label Informed Trajectory Prediction

Here, we introduce the concept of behavioral pseudo-labels for more accurate trajectory prediction.

1) Gumbel-Softmax Straight-Through Estimator

While the soft assignment represents good behavior clusters, such clusters are unsupervised and trained only on feature representations, meaning that they are still sub-optimal for any given task. This explains the sub-optimal prediction accuracy in existing methods [19], [64]. Here, we present a framework to improve the compatibility between the clusters and the task via fine-tuning the behavior latent.

To enable end-to-end fine-tuning of the behavior latent with a task objective, an operator is needed to connect the clustering and the prediction modules. We employ the Gumbel-Softmax straight-through estimator [74], which facilitates the gradient propagation and computes one-hot vectors representing the

pseudo-labels. The estimator uses a differentiable Softmax, as opposed to the non-differentiable Argmax, allowing end-to-end optimization. An agent's class label is l_j , where $j \in 1, \dots, k$ is the cluster center.

Apart from performance gains, as one-hot labels fit the human understanding of a class concept, they allow better interpretability via visualization tools. They are also immediately compatible with existing network architectures trained with ground-truth labels [26], [28], allowing effective adaptations.

2) Behavioral Pseudo-Label Informed SGCN

We adopt a Sparse Graph Convolution Network (SGCN) [10] as our backbone and introduce the pseudo-labels and a new loss function. SGCN has shown outstanding performance and is computationally efficient on pedestrian trajectory prediction [10]. It introduces sparsified spatial-temporal attention mechanism [11], [44], [75], which effectively models spatial interactions and temporal dependencies among agents. The sparse graph learning component removes spatial superfluous interactions and temporal motion tendencies, improving both computational speed and accuracy. In reality, our pseudo-label framework is compatible with a wide range of trajectory prediction networks.

We introduce the usage of semantic-goal features into SGCN, which enhances the prediction accuracy [22], [23]. To this end, we integrate the goal-retrieval operation [23] into the SGCN, we first subtract each observation step v_t in $t \in [1, T_{obs}]$ by the corresponding trajectory endpoint $v_{T_{pred}}$ as $\mathbf{v}_t = \mathbf{v}_t - \mathbf{v}_{T_{pred}}$. We then construct the spatial graph $\mathcal{G}_s = \{(\mathcal{V}_s, \mathcal{A}_s) | \mathcal{V}_s \in \mathbb{R}^{T_{obs} \times N \times D_s}, \mathcal{A}_s \in \mathbb{R}^{T_{obs} \times N \times N}\}$, where \mathcal{V}_s represents the spatial interactions among all agents at time step t , \mathcal{A}_s is the spatial adjacency matrix and D_s refers to the spatial feature dimension.

To add heterogeneity to the graph, we concatenate the pseudo-labels l to the trajectory feature vector for each agent at each time step as $\mathcal{V}_t^i = \text{concat}(v_t^i, l^i)$, $\forall t \in [1, T_{obs}]$ and $\forall i \in [1, N]$. Similarly, we establish the temporal graph $\mathcal{G}_t = \{(\mathcal{V}_t, \mathcal{A}_t) | \mathcal{V}_t \in \mathbb{R}^{N \times T_{obs} \times D_t}, \mathcal{A}_t \in \mathbb{R}^{N \times T_{obs} \times T_{obs}}\}$ to represent the temporal correlations of each individual agent during T_{obs} steps, where \mathcal{A}_t is the temporal adjacency matrix and D_t is the temporal feature dimension. Finally, these spatial and temporal goal-guided heterogeneous graphs are passed into SGCN for final trajectory prediction.

We propose a joint training strategy with a novel loss function to jointly optimize trajectory prediction and pseudo-label clustering. Thanks to our Gumbel-Softmax estimator, back-propagation is performed from the prediction all the way back to the VRNN encoder, resulting in better compatibility between the clustering and prediction modules. We present a combined loss:

$$\mathcal{L}_{final} = \mathcal{L}_{cluster} + \mathcal{L}_{prediction}, \quad (14)$$

where $\mathcal{L}_{cluster}$ is defined in Eq. 12, and $\mathcal{L}_{prediction}$ as:

$$\mathcal{L}_{prediction} = -\sum_{t=T_{obs}+1}^{T_{pred}} \log P(p_t | \hat{\mu}, \hat{\sigma}, \hat{\rho}), \quad (15)$$

where $\hat{\mu}$ and $\hat{\sigma}$ are the mean and variance of the bi-variate Gaussian distribution of trajectory prediction, and $\hat{\rho}$ represents the correlation coefficient.

IV. EXPERIMENTS

A. Datasets

We evaluate BP-SGCN on multiple benchmark datasets, including the Stanford Drone Dataset (SDD) [30], Argoverse 1 [34], ETH [20] and UCY [21], and the pedestrian-only version of SDD [35]. For pedestrian trajectory prediction, ETH/UCY consists of five pedestrian-only datasets (ETH, HOTEL, UNIV, ZARA1, ZARA2) with 1,536 pedestrians. Pedestrian-only SDD is the simplified version where non-pedestrian agents are removed. For heterogeneous trajectory prediction, we follow [28], [54], [76] that consider all trajectories, consisting of 8 scenes, 60 videos and 6 categories of traffic agents (i.e., pedestrians, bicyclists, skateboarders, carts, cars, and buses). Argoverse 1 consists of over 30K urban traffic scenarios that include 3 types of agents (i.e. AVs, agents, and others).

B. Experimental Setup

By default, we follow the experimental setup of [10], [13], using 3.2 seconds (8 frames) of observation trajectories to predict the next 4.8 seconds (12 frames). For pedestrian-only prediction, we employ the data augmentation approach introduced in [13] and the official leave-one-out strategy [37] during the training and validation. For heterogeneous trajectory prediction on Argoverse 1 dataset, our experimental setup and dataset split strategy follow [27], [53]. Specifically, we utilize 2 seconds (20 frames) of observation trajectories to predict the trajectories of all tracked objects over the subsequent 3 seconds (30 frames) within each scene. In particular, our experimental setup on the Argoverse 1 dataset for heterogeneous trajectory prediction predicts trajectories for all agents [27], [77], unlike methods focusing on a single agent [78] or two specific agents [49], our approach captures multi-agent interactions, reflecting real-world traffic complexity and improving predictive robustness, situational awareness, and adaptability to diverse urban environments.

During testing, we adhere to the standard protocol by generating 20 predictions for both heterogeneous [26], [28], [53] and pedestrian-only trajectory predictions [10], [37], [79]. This approach ensures our results are comparable to those established in the field. The sample with the lowest error is then used to compute the evaluation metrics. We employ the Average Displacement Error (ADE) and Final Displacement Error (FDE) [10], [12], [25], [37] as our evaluation metrics:

$$ADE = \frac{1}{(T_{pred} - T_{obs}) \times N} \sum_{i=1}^N \sum_{t=T_{obs}+1}^{T_{pred}} \|\hat{p}_t^i - p_t^i\|_2, \quad (16)$$

$$FDE = \frac{1}{N} \sum_{i=1}^N \|\hat{p}_t^i - p_t^i\|_2, t = T_{pred},$$

where \hat{p}_t^i represents the ground-truth trajectory coordinates.

TABLE I
RESULTS ON SDD FOR HETEROGENEOUS PREDICTION.

Methods	Venue	Year	GT Labels	SDD	
				ADE(\downarrow)	FDE(\downarrow)
Social-LSTM [25]	CVPR	2016	No	31.19	56.97
DESIRE [51]	CVPR	2017	No	19.25	34.05
MATF [80]	CVPR	2019	No	22.59	33.53
STGAT [11]	ICCV	2019	No	18.80	31.30
Multiverse [76]	CVPR	2020	No	14.78	27.09
SimAug [81]	ECCV	2020	No	10.27	19.71
NLNI [27]	ICCV	2021	Yes	15.90	26.30
STSF-Net [82]	TMM	2021	No	14.81	28.03
Semantic-STGCNN [26]	SMC	2021	Yes	18.12	29.70
V ² -Net [52]	ECCV	2022	No	<u>7.12</u>	<u>11.39</u>
Multiclass-SGCN [28]	ICIP	2022	Yes	14.36	25.99
TDOR [54]	CVPR	2022	No	8.60	13.90
CAPHA [83]	TVT	2023	No	9.13	14.34
VNAGT [84]	TVT	2023	Yes	9.67	17.22
SFEM-GCN [29]	TIV	2024	Yes	15.31	25.72
SMGCN [58]	IJCAI	2024	Yes	20.89	36.84
BP-SGCN (Ours)			No	6.94	9.57

TABLE II
RESULTS ON ARGOVERSE I FOR HETEROGENEOUS PREDICTION.

Methods	Venue	Year	GT Labels	Argoverse I	
				ADE(\downarrow)	FDE(\downarrow)
Social-LSTM [25]	CVPR	2016	No	1.39	2.57
DESIRE [51]	CVPR	2017	No	0.90	1.45
R2P2-MA [85]	ECCV	2018	No	1.11	1.77
MATFG [80]	CVPR	2019	No	1.26	2.31
CAM [86]	ECCV	2020	No	1.13	2.50
MFP [87]	NeurIPS	2020	No	1.40	2.68
Social-STGCNN [12]	CVPR	2020	No	1.31	2.34
NLNI [27]	ICCV	2021	Yes	0.79	1.26
DD [88]	Inf. Sci.	2022	No	<u>0.74</u>	1.28
HRG+HSG [53]	TITS	2023	No	0.85	1.12
BIP-Tree [77]	TITS	2023	No	0.78	1.35
BP-SGCN (Ours)			No	0.69	<u>1.15</u>

C. Quantitative Evaluation

1) Heterogeneous Prediction

TABLE I compares our BP-SGCN with previous state-of-the-art methods on heterogeneous SDD. These methods can be categorized into three groups based on the input features, including trajectory-only [11], [25], [29], [63], [82], trajectory with ground-truth labels [26]–[29], [58], [84], and trajectory with extra scene features such as scene semantics [51], [52], [54], [76], [81], [83]. BP-SGCN outperforms all the methods that utilize ground-truth agent class labels [26]–[29], [58], [84]. Compared to the best method VNAGT [84], BP-SGCN demonstrates the superiority by reducing ADE/FDE by 28.23%/44.43%. Crucially, for SOTA approaches that incorporate scene semantic features such as V²-Net [52] and TDOR [54], our BP-SGCN improves the performance by reducing ADE/FDE by 2.5%/15.9% compared to V²-Net and 19.3%/31.2% compared to TDOR. The results indicate that without the need for additional inputs, our BP-SGCN can still achieve SOTA performance in heterogeneous trajectory prediction.

TABLE II compares the BP-SGCN with those state-of-the-art methods in heterogeneous trajectory prediction on Argoverse I, following the setup in [27], [53], [77]. Results show that our

BP-SGCN outperforms all the methods by a significant margin, especially in the ADE metric. BP-SGCN surpasses NLNI [27], which integrates ground-truth labels, by reducing 12.7% in ADE and 8.7% in FDE, further showcasing the effectiveness of our proposed pseudo-label module. Notably, although NLNI utilizes label-based category features, its performance is limited by the simplistic nature of the “GT Labels” in the Argoverse 1 dataset, which are broadly classified as “1 AV” (1 Autonomous Vehicle), “1 Focal” (the primary vehicle whose trajectory is predicted), and “N other” (other tracked objects, which can include vehicles, pedestrians, or bicycles). This coarse categorization restricts the algorithm’s ability to accurately capture and analyze the nuanced interactions among diverse traffic agents. In contrast, BP-SGCN effectively overcomes these constraints by conducting a comprehensive analysis of the behavior dynamics of all agents within the scene. By employing our advanced pseudo-label module, we significantly enhance the representational capabilities of our system, leading to markedly improved prediction accuracy across diverse traffic scenarios. This improvement is achieved without the need for direct matching with ground-truth labels, demonstrating the robustness and adaptability of our approach in interpreting complex interactive behaviors. Importantly, DD [88] and HRG+HSG [53] achieve comparable performance on ADE and FDE mainly due to the use of scene images that better capture the interactions between traffic agents and environments, our BP-SGCN still shows the best ADE performance compared to these methods.

2) Pedestrian Prediction

For ETH/UCY, we conduct quantitative comparisons with a wide range of methods with various techniques, as shown in TABLE III. Following [41], [92], we compare with methods utilizing trajectory data only.

For distribution-based methods, Social-LSTM [25] introduces bi-variate Gaussian distribution to sample predictions from the trained mean and variance, which is widely used in recently published methods [10], [12], [23], [63], [97], [106]. Following this, our BP-SGCN also uses the bi-variate Gaussian distribution to represent the distribution parameters of the predicted trajectories. It outperforms almost all methods under this setting by a significant margin. In addition, both ExpertTraj [23] and our BP-SGCN utilize the goal-retrieve mechanism but we have a significant improvement of 10.5% in ADE and 33.3% in FDE.

For generative-based methods, Social-GAN [37] is the pioneer method that introduces GANs [38] to generate trajectories with special pooling modules. PECNet [24] utilizes the CVAEs [50] to generate trajectories conditioned on the pre-sampled goal points, which add an extra constraint to the predicted trajectories for better accuracy. Methods like [45]–[47], [91], [93], [101] follow the CVAEs basis to train the encoder with ground-truth trajectories for better latent representations. MID [40] and LED [41] further introduce the diffusion models [39] to enhance training and reduce mode collapses. Results reveal that our BP-SGCN outperforms generative-based methods.

TABLE III
RESULTS ON ETH/UCY ON PEDESTRIAN-ONLY PREDICTION;
- DENOTES MISSING RESULT DUE TO UNAVAILABILITY FROM ORIGINAL AUTHORS.

Method	Venue	Year	ETH	HOTEL	UNIV	ZARA1	ZARA2	AVG
			ADE(↓)/FDE(↓)	ADE(↓)/FDE(↓)	ADE(↓)/FDE(↓)	ADE(↓)/FDE(↓)	ADE(↓)/FDE(↓)	ADE(↓)/FDE(↓)
Social LSTM [25]	CVPR	2016	1.09/2.35	0.79/1.76	0.67/1.40	0.47/1.00	0.56/1.17	0.72/1.54
Social GAN [37]	CVPR	2018	0.87/1.62	0.67/1.37	0.76/1.52	0.35/0.68	0.42/0.84	0.61/1.21
Social-STGCNN [12]	CVPR	2020	0.64/1.11	0.49/0.85	0.44/0.79	0.34/0.53	0.30/0.48	0.44/0.75
PECNet [24]	ECCV	2020	0.54/0.87	0.18/0.24	0.35/0.60	0.22/0.39	0.17/0.30	0.29/0.48
SGCN [10]	CVPR	2021	0.63/1.03	0.32/0.55	0.37/0.70	0.29/0.53	0.25/0.45	0.37/0.65
AgentFormer [45]	ICCV	2021	0.45/0.75	0.14/0.22	0.25/0.45	0.18/0.30	0.14/0.24	0.23/0.39
PCCSNet [19]	ICCV	2021	0.28/0.54	0.11/0.19	0.29/0.60	0.21/0.44	0.15/0.34	0.21/0.42
ExpertTraj [23]	ICCV	2021	0.37/0.65	0.11/0.15	0.20/0.44	0.15/0.31	0.12/0.25	0.19/0.36
STSF-Net [82]	TMM	2021	0.63/1.13	0.24/0.43	0.28/0.52	0.23/0.45	0.21/0.41	0.32/0.59
Social-Implicit [13]	ECCV	2022	0.66/1.44	0.20/0.36	0.31/0.60	0.25/0.50	0.22/0.43	0.33/0.67
GP-Graph [63]	ECCV	2022	0.43/0.63	0.18/0.30	0.24/0.42	0.17/0.31	0.15/0.29	0.23/0.39
Social-VAE [47]	ECCV	2022	0.41/0.58	0.13/0.19	0.21/0.36	0.17/0.29	0.13/0.22	0.21/0.33
MemoNet [48]	CVPR	2022	0.40/0.61	0.11/0.17	0.24/0.43	0.18/0.32	0.14/0.24	0.21/0.35
GroupNet [46]	CVPR	2022	0.46/0.73	0.15/0.25	0.26/0.49	0.21/0.39	0.17/0.33	0.25/0.44
MID [40]	CVPR	2022	0.39/0.66	0.13/0.22	0.22/0.45	0.17/0.30	0.13/0.27	0.21/0.38
GTPO [89]	TNNLS	2022	0.63/0.98	0.19/0.30	0.35/0.60	0.20/0.32	0.18/0.31	0.31/0.50
Graph-TERN [90]	AAAI	2023	0.42/0.58	0.14/0.23	0.26/0.45	0.21/0.37	0.17/0.29	0.24/0.88
MSRL [91]	AAAI	2023	0.28/0.47	0.14/0.22	0.24/0.43	0.17/0.30	0.14/0.23	0.19/0.33
LED [41]	CVPR	2023	0.39/0.58	0.11/0.17	0.26/0.43	0.18/0.26	0.13/0.22	0.21/0.33
EqMotion [92]	CVPR	2023	0.40/0.61	0.12/0.18	0.23/0.43	0.18/0.32	0.13/0.23	0.21/0.35
FEND [64]	CVPR	2023	-	-	-	-	-	0.17/0.32
EigenTrajectory [93]	ICCV	2023	0.36/0.53	0.12/0.19	0.24/0.43	0.19/0.33	0.14/0.24	0.21/0.34
TUTR [94]	ICCV	2023	0.40/0.61	0.11/0.18	0.23/0.42	0.18/0.34	0.13/0.25	0.21/0.36
SICNet [95]	ICCV	2023	0.27/0.45	0.11/0.16	0.26/0.46	0.19/0.33	0.13/0.26	0.19/0.33
TP-EGT [96]	TITS	2023	0.41/0.68	0.13/0.21	0.29/0.50	0.18/0.30	0.16/0.27	0.23/0.39
DynGroupNet [97]	NN	2023	0.42/0.66	0.13/0.20	0.24/0.44	0.19/0.34	0.15/0.28	0.23/0.38
SMEMO [98]	TPAMI	2024	0.39/0.59	0.14/0.20	0.23/0.41	0.19/0.32	0.15/0.26	0.22/0.35
STGlow [99]	TNNLS	2024	0.31/0.49	0.09/0.14	0.16/0.33	0.12/0.24	0.09/0.19	0.15/0.28
MRGTraj [100]	TCSVT	2024	0.28/0.47	0.21/0.39	0.33/0.60	0.24/0.44	0.22/0.41	0.26/0.46
HighGraph [101]	CVPR	2024	0.40/0.55	0.13/0.17	0.20/0.33	0.17/0.27	0.11/0.21	0.20/0.30
PPT [102]	ECCV	2024	0.36/0.51	0.11/0.15	0.22/0.40	0.17/0.30	0.12/0.21	0.20/0.31
BP-SGCN (Ours)			0.33/0.47	0.10/0.14	0.17/0.26	0.13/0.19	0.10/0.16	0.17/0.24

¹ For ExpertTraj [23], the discrepancy from the original paper arises due to an error highlighted by the authors: https://github.com/JoeHEZHAO/expert_traj

TABLE IV
RESULTS ON THE PEDESTRIAN-ONLY VERSION OF SDD.

Methods	Venue	Year	SDD-human	
			ADE(↓)	FDE(↓)
STGAT [11]	ICCV	2019	0.58	1.11
Social-Ways [103]	CVPRW	2019	0.62	1.16
DAG-Net [79]	ICPR	2020	0.53	1.04
Social-implicit [13]	ECCV	2022	0.47	0.89
WTGCN [104]	IJMLC	2024	0.43	0.72
IGGCN [105]	DSP	2024	0.44	0.71
BP-SGCN (Ours)			0.28	0.41

For transformer-based methods, TUTR [94] proposes a novel global prediction system incorporated with a motion-level transformer encoder and a social-level transformer decoder for accurate trajectory representation. MRGTraj [100] introduces a non-autoregressive enhanced transformer decoder for trajectory prediction. PPT [102] proposes multi-stage transformer progressively modeling trajectories. STGlow [99] further introduces the flow-based generative framework with dual-graphormer to precisely model motion distributions. Compared to STGlow, our BP-SGCN achieves comparable ADE with STGlow, while reducing FDE by 14%.

Besides these categories, LSTM decoder-based methods [19], [82], [89], [96] and [64] directly predict trajectories using LSTM decoder, which also show comparable results to the

transformer-based methods. Social-implicit [13] introduces the concept of implicit maximum likelihood estimation mechanism. Memonet [48], SICNet [95] and SMEMO [98] incorporate memory bank/module concepts into the system, demonstrating considerable performance. Notably, SICNet presents the best results on ETH subset in both ADE and FDE metrics compared with all other methods. Graph-TERN [90] shows a novel trajectory refinement module that first samples the endpoint and then linearly interpolates the predictions. EqMotion [92] further introduces the concepts of invariance and equivariance into trajectory prediction to learn motion patterns. Nevertheless, results in TABLE III illustrate that our BP-SGCN outperforms all of these methods.

For pedestrian-only SDD, TABLE IV highlights the comparative performance of our BP-SGCN, which secures substantial improvements over all listed models, including the latest STOA models, WTGCN [104] and IGGCN [105]. Specifically, BP-SGCN achieves a 35% reduction in ADE compared to WTGCN and a 42% reduction in FDE compared to IGGCN. The results in both heterogeneous SDD and pedestrian-only SDD show the superiority of our BP-SGCN in multiple scenarios.

D. Qualitative Evaluation

Fig. 4 shows the t-SNE [65] visualization of the pseudo-labels in the latent space during unsupervised deep clustering




Fig. 4. The t-SNE visualization of pseudo-class clustering on SDD ($k=6$) during unsupervised deep clustering. (a) 0 epochs (initialized by k-means), (b) 200 epochs, (c) 800 epochs.

training. At epoch 0, k-means initializes the cluster centers. As the latent is not optimized, class boundaries are unclear. After the training, the VRNN encoder generates a more representative latent for better clustering of the agents’ behaviors.

Fig. 5 and Fig. 6 visualize the trajectory predictions for the SDD and ETH/UCY datasets, respectively. Blue and red dots represent observed and ground-truth future trajectories, respectively. For the SDD dataset, we visualize the predictions in Fig. 5, where light blue indicates the predicted distributions and yellow dots represent the predicted single trajectory. The visualizations demonstrate that our BP-SGCN exhibits superior performance compared to methods integrating ground-truth labels [26], [28] in three challenging scenarios characterized by complex social interactions among agents.

In Fig. 6, we visualize the predicted distribution in the ETH/UCY datasets across various scenarios, encompassing both simple and complex interactions, and compare our method with SGCN [10] and GP-Graph [63]. We visualize the parameterized distribution of future trajectories, as they are the learning objective of these methods. Qualitative comparisons reveal that our predicted distributions closely align with the ground truth and adeptly capture the non-linear trajectories. Specifically, scenario (a) illustrates a scene with numerous pedestrians on the street engaging in complex interactions, such as meeting, colliding, and standing still. While all the predicted distributions can accurately represent linear trajectories, both SGCN and GP-Graph falter in predicting the movements of pedestrians exhibiting non-linear behaviors. In contrast, BP-SGCN consistently generates plausible predictions. Scenario (b) displays four stationary pedestrians; however, both SGCN and GP-Graph yield wrong predictions, whereas BP-SGCN accurately captures the static behaviors. In scenario (c), the predicted distributions from both SGCN and GP-Graph demonstrate significant overlaps, leading to a heightened risk of predicted collisions. On the other hand, BP-SGCN’s predictions show reduced overlaps. In scenario (d), while GP-Graph continues to display overlap issues, SGCN exhibits overconfidence in its predictions, resulting in a lack of diversity and a propensity to deviate from the ground truth. BP-SGCN effectively addresses both of these challenges, striking a balance between prediction accuracy and diversity.




Fig. 5. Visualization of trajectory prediction on SDD of Semantic-STGCNN [26], Multiclass-SGCN [28], and BP-SGCN (ours). Blue and red represent observed and ground-truth trajectories respectively, yellow represents the predicted trajectory and light-blue shade represents the predicted distribution.

E. Ablation Study and Parameter Analysis

1) Cluster Number Analysis

The effects of cluster number on heterogeneous datasets are shown in TABLE V (Heterogeneous SDD) and TABLE VI (Argoverse 1). The results on pedestrian-only datasets are shown in TABLE VII (ETH/UCY) and TABLE VIII (Pedestrian-only SDD). In general, the cluster number depends on the diversity of behaviors, which is strongly correlated with the location. For instance, choosing six clusters for SDD is reasonable given the presence of six types of agents, and this choice yields good performance. Tuning the cluster number for a scene provides extra improvements, and this only has to be done once. These results further reflect that the heterogeneous dataset is more sensitive to the cluster numbers and the pedestrian dataset results exhibit diminished sensitivity, attributable to the inherent behavioral homogeneity and comparatively lower variance observed in human actions. Note that due to its large data size, for Argoverse 1, we run ablation studies and parameter analysis using a partial dataset in a simplified setup.

Notably, in our experiments on cluster numbers, a cluster’s number equal to 1 denotes that there is no pseudo-label applied on each agent because all the agents are considered to belong to the same class, and consequently the model performance relies solely on the trajectories themselves. In particular, results in TABLE VII and TABLE VIII demonstrate that, within datasets exclusively comprising pedestrian agents, our BP-SGCN model is adept at discerning the nuanced variances in their movement patterns. Despite the apparent homogeneity of the agents as pedestrians, our analysis reveals intrinsic behavioral differentiations that our model capitalizes on to significantly improve prediction accuracy. This not only underscores the importance of individualized learning even

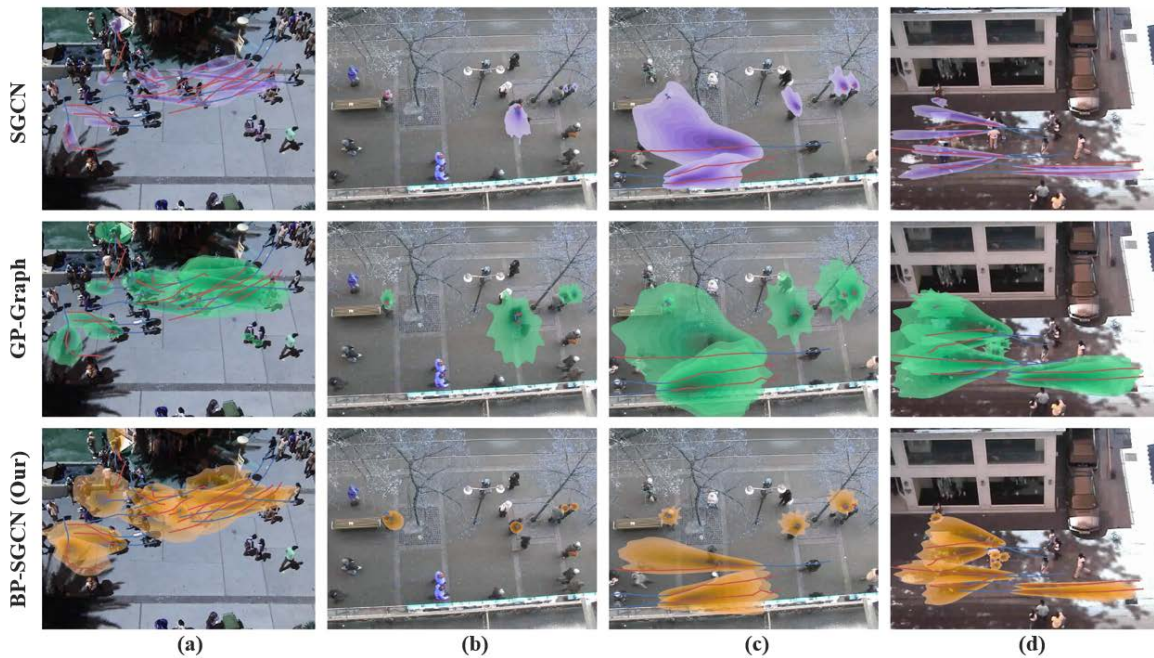


Fig. 6. Visualization of the trajectory prediction on ETH/UCY in the scenario of pedestrian walking behaviors. Past trajectories are shown in blue, and ground-truth trajectories are in red. (a) shows the pedestrians in a crowded scenario with complex interactions. (b) shows the scene where four pedestrians are almost static. (c) and (d) show scenes including multiple pedestrian behaviors, such as walking, meeting, and standing.

TABLE V
CLUSTER NUMBER ANALYSIS ON HETEROGENEOUS SDD.

Clusters	ADE(↓)	FDE(↓)
1	7.26	10.03
3	7.11	9.81
6	6.94	9.57
9	7.03	9.74
12	7.58	10.92

TABLE VI
CLUSTER NUMBER ANALYSIS ON ARGVERSE 1.

Clusters	ADE(↓)	FDE(↓)
1	0.86	1.63
3	0.80	1.45
6	0.69	1.15
9	0.79	1.47

among seemingly similar entities, but also showcases the efficacy of our model in enhancing predictive outcomes by leveraging these subtle distinctions.

2) Network Components Analysis

TABLE IX shows ablation studies to evidence the effectiveness of network components used in BP-SGCN on heterogeneous and pedestrian-only SDD. The “No Deep Clustering” setup uses k-means cluster centers directly for trajectory prediction, and therefore does not implement unsupervised deep learning and end-to-end fine-tuning. The “No Gumbel-Softmax” setup directly concatenates the soft assignment to the trajectory features for trajectory prediction. The “No End-to-End Training” setup uses only $L_{prediction}$ to optimize the trajectory prediction

TABLE VII
CLUSTER NUMBER ANALYSIS ON ETH/UCY.

Clusters	ADE(↓) / FDE(↓)				
	ETH	HOTEL	UNIV	ZARA1	ZARA2
1	0.37/0.51	0.14/0.19	0.27/0.37	0.15/0.21	0.24/0.34
2	0.37/0.52	0.15/0.21	0.18/0.27	0.20/0.37	0.25/0.35
3	0.45/0.61	0.10/0.14	0.27/0.36	0.24/0.33	0.17/0.34
4	0.33/0.47	0.12/0.16	0.27/0.37	0.14/0.20	0.10/0.16
5	0.36/0.50	0.17/0.22	0.18/0.27	0.13/0.19	0.12/0.18
6	0.39/0.53	0.15/0.21	0.18/0.27	0.15/0.21	0.11/0.17
7	0.37/0.51	0.11/0.15	0.17/0.26	0.26/0.37	0.13/0.19

TABLE VIII
CLUSTER NUMBER ANALYSIS ON PEDESTRIAN-ONLY SDD.

Clusters	ADE(↓)	FDE(↓)
1	0.33	0.49
3	0.28	0.41
6	0.47	0.72
9	0.31	0.47

module but not the deep clustering module; here, the Gumbel-Softmax estimator is substituted with the non-differentiable Argmax function. Results from both the heterogeneous and pedestrian datasets emphasize the significance of all the proposed components in BP-SGCN.

In addition, our proposed Goal-Guided SGCN module utilizes the spatial attention and temporal attention mechanism to enhance the final prediction accuracy. We conduct experiments on ETH/UCY datasets to validate the effectiveness of these two modules. The results shown in TABLE X indicate that both spatial attention and temporal attention modules are important for the best performance.

TABLE IX
NETWORK COMPONENTS ANALYSIS ON HETEROGENEOUS SDD (UPPER)
AND PEDESTRIAN-ONLY SDD (LOWER).

Method	ADE(↓)	FDE(↓)
BP-SGCN (No Deep Clustering)	7.52	10.50
BP-SGCN (No Gumbel-Softmax)	7.65	10.85
BP-SGCN (No End-to-End Training)	10.82	15.32
BP-SGCN (Ours)	6.94	9.57

Method	ADE(↓)	FDE(↓)
BP-SGCN (No Deep Clustering)	0.30	0.44
BP-SGCN (No Gumbel-Softmax)	0.40	0.60
BP-SGCN (No End-to-End Training)	0.30	0.46
BP-SGCN (Ours)	0.28	0.41

TABLE X
PREDICTION MODULE ANALYSIS ON ETH/UCY DATASETS.

Method	ADE(↓)	FDE(↓)
BP-SGCN (No Spatial Attention)	0.25	0.30
BP-SGCN (No Temporal Attention)	0.28	0.35
BP-SGCN (Ours)	0.17	0.24

3) Trajectory Prediction Loss Analysis

As discussed above, we propose a cascaded training strategy with a novel loss function to jointly optimize trajectory prediction and pseudo-label clustering, defined as:

$$\mathcal{L}_{final} = \mathcal{L}_{prediction} + \mathcal{L}_{cluster}. \quad (17)$$

In the proposed loss function, $\mathcal{L}_{prediction}$ and $\mathcal{L}_{cluster}$ contribute equally to the final loss \mathcal{L}_{final} . We conduct an ablation study by introducing a weighted sum of losses with a new hyperparameter λ to explore the effect and contribution of the two losses on trajectory prediction on both heterogeneous and pedestrian-only SDD datasets:

$$\mathcal{L}_{final} = \lambda\mathcal{L}_{prediction} + (1 - \lambda)\mathcal{L}_{cluster}. \quad (18)$$

Here, we analyze the effect of λ . For the proposed BP-SGCN, the default value of λ can be considered as 0.5, as both losses contribute equally to the final loss. We further adjust the value of λ as 0.25, and 0.75, respectively. The experimental results presented in TABLE XI show that the performance of BP-SGCN reaches its peak when the ratio of $\mathcal{L}_{prediction}$ and $\mathcal{L}_{cluster}$ is equal, as presented in the main paper, which further indicates that the trajectory prediction and pseudo-label clustering modules are equally important for the overall trajectory prediction performance.

4) Clustering Features Analysis

Finally, TABLE XII shows ablation studies on heterogeneous and pedestrian-only SDD datasets with regard to the geometric features used for behavior clustering. These features play a pivotal role, enabling our unsupervised deep clustering module to differentiate agent behaviors effectively. The outcomes highlight the outstanding performance of our proposed features, which integrate relative angle and acceleration magnitude.

TABLE XI
LOSS WEIGHT ANALYSIS BETWEEN $\mathcal{L}_{prediction}$ AND $\mathcal{L}_{cluster}$ ON
HETEROGENEOUS SDD (UPPER) AND PEDESTRIAN-ONLY SDD (LOWER).

Method	ADE(↓)	FDE(↓)
BP-SGCN ($\lambda = 0.25$)	19.33	24.26
BP-SGCN ($\lambda = 0.75$)	7.08	9.84
BP-SGCN (Ours)	6.94	9.57

Method	ADE(↓)	FDE(↓)
BP-SGCN ($\lambda = 0.25$)	0.46	0.70
BP-SGCN ($\lambda = 0.75$)	0.31	0.46
BP-SGCN (Ours)	0.28	0.41

F. Model Complexity and Inference Time Analysis

To verify the efficiency of our proposed method, we conduct experiments on inference time and model parameters with existing mainstream trajectory prediction frameworks. As demonstrated in TABLE XIII, our method is inferior to EigenTrajectory [93] and better than all other methods in terms of inference time and model parameters. We leave it as future work to improve the efficiency of our BP-SGCN with more advanced sequential modeling methods such as Transformers [44] and State Space Models (SSMs) [107], [108].

TABLE XII
CLUSTERING FEATURES ANALYSIS ON HETEROGENEOUS SDD (UPPER)
AND PEDESTRIAN-ONLY SDD (LOWER).

Method	ADE(↓)	FDE(↓)
BP-SGCN (Relative Angle)	19.52	34.05
BP-SGCN (Acceleration Magnitude)	9.07	13.02
BP-SGCN (Ours)	6.94	9.57

Method	ADE(↓)	FDE(↓)
BP-SGCN (Relative Angle)	0.45	0.68
BP-SGCN (Acceleration Magnitude)	0.42	0.63
BP-SGCN (Ours)	0.28	0.41

TABLE XIII
COMPARISON OF THE PROPOSED APPROACHES IN TERMS OF
NUMBER OF PARAMETER AND INFERENCE TIME.

Methods	Venue	Year	Param $\times 10^6$	Infer. Time/Iter.
ExpertTraj [23]	ICCV	2021	0.32	130 ms
Social-VAE [47]	ECCV	2022	5.69	1110 ms
GroupNet [46]	CVPR	2022	3.14	-
MSRL [91]	AAAI	2023	11.32	970 ms
EqMotion [92]	CVPR	2023	2.08	800 ms
TUTR [94]	ICCV	2023	0.44	360 ms
EigenTrajectory [93]	ICCV	2023	0.02	72 ms
BP-SGCN (Ours)			0.13	110 ms

Moreover, we validate the stability and reliability of our BP-SGCN on heterogeneous trajectory prediction by 10 experiments. Results shown in TABLE XIV showcase the stability of our method.

G. Discussion

In our experiments, we observed that methods [47], [91], [93], [94] tailored exclusively for pedestrians exhibit a sensitivity to




Fig. 7. Visualization of the trajectory prediction of BP-SGCN in different social scenarios including positive predictions and negative predictions (we highlight erroneous predictions inside the white boxes). Past trajectories are shown in blue, ground-truth trajectories are in red, predicted trajectories are shown in yellow, and distributions are shown in light blue.

TABLE XIV
STABILITY TESTS ON ARGOVERSE 1 AND HETEROGENEOUS VERSION OF SDD

Methods	Argoverse 1		SDD	
	ADE(↓)	FDE(↓)	ADE(↓)	FDE(↓)
BP-SGCN (ours)	0.68 ± 0.031	1.16 ± 0.034	6.97 ± 0.069	9.59 ± 0.043

the threshold settings that dictate the count of nearby agents. These methods, while ensuring state-of-the-art performance in pedestrian-only trajectory prediction, perform sub-optimally in heterogeneous scenarios due to the challenge of predefining neighbors. The result is shown in TABLE XV. Unlike these pedestrian-specific approaches, which require manual neighbor selection based on metrics like relative distances, our BP-SGCN model automatically considers all proximate agents as initial neighbors, adaptively filtering out the less relevant ones. Thus, our proposed BP-SGCN is better than these methods in heterogeneous trajectory prediction.

TABLE XV
RESULTS BY PEDESTRIAN-ONLY METHODS ON THE HETEROGENEOUS VERSION OF SDD.

Methods	Venue	Year	SDD	
			ADE(↓)	FDE(↓)
Social-VAE + FPC [47]	ECCV	2022	9.41	13.49
MSRL [91]	AAAI	2023	10.72	16.15
EigenTrajectory [93]	ICCV	2023	8.85	15.15
TUTR [94]	ECCV	2023	8.93	15.66
BP-SGCN (Ours)			6.94	9.57

Next, we showcase inaccurate predictions made by our BP-SGCN and delve into the method’s limitations. As depicted in Fig. 7, the first row illustrates the BP-SGCN’s proficiency in accurately predicting trajectories across various social contexts. Nonetheless, the second row highlights instances where our BP-SGCN falls short, particularly in scenarios where: 1) trajectories undergo abrupt changes; 2) paths are highly erratic and frequently alter; and 3) social dynamics become exceedingly intricate with numerous agents involved. Looking ahead, our

objective is to rectify these inaccuracies by enhancing BP-SGCN’s capabilities through the incorporation of cutting-edge deep learning methodologies, including Transformers [44] and Diffusion models [41], among others.

The quantity of behavior clusters is an adjustable hyperparameter. We manually select the number of clusters for the unsupervised deep clustering module. This approach brings several challenges, including subjectivity and potential bias, scalability issues, and potential impacts on model performance due to overfitting or underfitting. Moreover, the optimal number of clusters is sensitive to the datasets, which further complicates the selection process. Especially in heterogeneous scenarios, the high variance between different types of agents’ motions makes it challenging to identify the best number of clusters to represent behavior features accurately than pedestrian-only scenarios. In the future, we aim to scrutinize the behavior distributions of traffic agents more closely and dynamically estimate the optimal number of clusters [109], [110].

Despite BP-SGCN’s effectiveness in both heterogeneous and pedestrian-only trajectory prediction, another notable limitation of our model is its current omission of scene semantic features. Although only using trajectories as inputs brings the benefit of computation efficiency and emphasizes the importance of behavior motions, the integration of agent interactions with their surrounding environment can benefit in developing effective trajectory prediction models for use in real-world scenarios [22], [52], [54], [111]. Recognizing this, a significant direction for our future is to explore how to effectively combine trajectory data with scene semantic features to capture the interactions between static barriers and dynamic agents. We hypothesize that this will not only enhance the model’s prediction accuracy, but also improve the refinement of pseudo-label identification by leveraging the rich context provided by environmental cues.

V. CONCLUSION

In this work, we introduce BP-SGCN for heterogeneous and pedestrian trajectory prediction, showcasing its superior performance compared to existing models. In particular, we

introduce the concept of behavioral pseudo-labels, which effectively represent the different behavior clusters of agents and do not require extra ground-truth information. BP-SGCN includes a deep unsupervised clustering module that learns the pseudo-label, as well as a pseudo-label informed sparse graph convolution network for trajectory prediction. It implements a cascaded training scheme that first learns the pseudo-labels in an unsupervised manner, and then fine-tunes the labels by optimizing the network end-to-end for better compatibility.

Beyond pedestrian scenarios, BP-SGCN also shows promising potential in broader domains. In robotic path planning [1], [99], BP-SGCN can enhance collision avoidance systems through behavioral pattern analysis [60], [89] of surrounding agents, facilitating more effective navigation in intricate settings. Additionally, in video monitoring and surveillance systems as suggested in [10], [99], BP-SGCN can enhance anomaly detection through behavioral pattern analysis of system dynamics, enabling early detection of potential operational irregularities. These applications demonstrate the applicability of BP-SGCN in modeling interactive behaviors across different domains, highlighting its potential for various real-world trajectory prediction tasks.

ACKNOWLEDGEMENT


This research is supported in part by the EPSRC NorthFutures project (ref: EP/X031012/1).

REFERENCES

- [1] Y. Luo, P. Cai, A. Bera, D. Hsu, W. S. Lee, and D. Manocha, "Porca: Modeling and planning for autonomous driving among many pedestrians," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3418–3425, 2018.
- [2] W. Yang, S. Li, and X. Luo, "Data driven vibration control: A review," *IEEE/CAA J. Autom. Sin.*, vol. 11, no. 9, pp. 1898–1917, 2024.
- [3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 779–788.
- [4] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. Int. Conf. Learn. Represent.*, 2017.
- [5] M. Li, A. Micheli, Y. G. Wang, S. Pan, P. Lió, G. S. Gnecco, and M. Sanguineti, "Guest editorial: Deep neural networks for graphs: Theory, models, algorithms, and applications," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4367–4372, 2024.
- [6] Y. Zheng, M. Jin, S. Pan, Y.-F. Li, H. Peng, M. Li, and Z. Li, "Toward graph self-supervised learning with contrastive adjusted zooming," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 7, pp. 8882–8896, 2024.
- [7] J. Li, R. Zheng, H. Feng, M. Li, and X. Zhuang, "Permutation equivariant graph framelets for heterophilous graph learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 9, pp. 11 634–11 648, 2024.
- [8] D. Wu, Y. He, and X. Luo, "A graph-incorporated latent factor analysis model for high-dimensional and sparse data," *IEEE Trans. Emerg. Top. Comput.*, vol. 11, no. 4, pp. 907–917, 2023.
- [9] X. Luo, H. Wu, Z. Wang, J. Wang, and D. Meng, "A novel approach to large-scale dynamically weighted directed network representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 12, pp. 9756–9773, 2022.
- [10] L. Shi, L. Wang, C. Long, S. Zhou, M. Zhou, Z. Niu, and G. Hua, "Sgcn: Sparse graph convolution network for pedestrian trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 8994–9003.
- [11] Y. Huang, H. Bi, Z. Li, T. Mao, and Z. Wang, "Stgat: Modeling spatial-temporal interactions for human trajectory prediction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 6272–6281.
- [12] A. Mohamed, K. Qian, M. Elhoseiny, and C. Claudel, "Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 14 424–14 432.
- [13] A. Mohamed, D. Zhu, W. Vu, M. Elhoseiny, and C. Claudel, "Social-implicit: Rethinking trajectory prediction evaluation and the effectiveness of implicit maximum likelihood estimation," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 463–479.
- [14] B. Xu, X. Shu, J. Zhang, G. Dai, and Y. Song, "Spatiotemporal decouple-and-squeeze contrastive learning for semisupervised skeleton-based action recognition," *IEEE Trans. Neural Netw. Learn. Syst.*, 2023.
- [15] X. Shu, B. Xu, L. Zhang, and J. Tang, "Multi-granularity anchor-contrastive representation learning for semi-supervised skeleton-based action recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022.
- [16] B. Xu, X. Shu, and Y. Song, "X-invariant contrastive augmentation and representation learning for semi-supervised skeleton-based action recognition," *IEEE Trans. Image Process.*, vol. 31, pp. 3852–3867, 2022.
- [17] B. Xu and X. Shu, "Pyramid self-attention polymerization learning for semi-supervised skeleton-based action recognition," *arXiv preprint arXiv:2302.02327*, 2023.
- [18] T. Qiao, Q. Men, F. W. B. Li, Y. Kubotani, S. Morishima, and H. P. H. Shum, "Geometric features informed multi-person human-object interaction recognition in videos," in *Proc. Eur. Conf. Comput. Vis.*, 2022.
- [19] J. Sun, Y. Li, H.-S. Fang, and C. Lu, "Three steps to multimodal trajectory prediction: Modality clustering, classification and synthesis," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 13 250–13 259.
- [20] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.* IEEE, 2009, pp. 261–268.
- [21] A. Lerner, Y. Chrysanthou, and D. Lischinski, "Crowds by example," in *Computer graphics forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 655–664.
- [22] K. Mangalam, Y. An, H. Girase, and J. Malik, "From goals, waypoints & paths to long term human trajectory forecasting," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 15 233–15 242.
- [23] H. Zhao and R. P. Wildes, "Where are you heading? dynamic trajectory prediction with expert goal examples," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 7629–7638.
- [24] K. Mangalam, H. Girase, S. Agarwal, K.-H. Lee, E. Adeli, J. Malik, and A. Gaidon, "It is not the journey but the destination: Endpoint conditioned trajectory prediction," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 759–776.
- [25] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 961–971.
- [26] B. A. Rainbow, Q. Men, and H. P. Shum, "Semantics-stgcnn: A semantics-guided spatial-temporal graph convolutional network for multi-class trajectory prediction," in *IEEE Int. Conf. Syst. Man Cybern.* IEEE, 2021, pp. 2959–2966.
- [27] F. Zheng, L. Wang, S. Zhou, W. Tang, Z. Niu, N. Zheng, and G. Hua, "Unlimited neighborhood interaction for heterogeneous trajectory prediction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 13 168–13 177.
- [28] R. Li, S. Katsigiannis, and H. P. Shum, "Multiclass-sgcn: Sparse graph-based trajectory prediction with agent class embedding," in *IEEE Int. Conf. Image Process.* IEEE, 2022, pp. 2346–2350.
- [29] Q. Du, X. Wang, S. Yin, L. Li, and H. Ning, "Social force embedded mixed graph convolutional network for multi-class trajectory prediction," *IEEE Trans. Intell. Veh.*, pp. 1–11, 2024.
- [30] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 549–565.
- [31] J. Xie, R. Girshick, and A. Farhadi, "Unsupervised deep embedding for clustering analysis," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 478–487.
- [32] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, "A recurrent latent variable model for sequential data," *Proc. Adv. Neu. Inf. Process. Syst.*, vol. 28, 2015.
- [33] W. Yang, S. Li, Z. Li, and X. Luo, "Highly accurate manipulator calibration via extended kalman filter-incorporated residual neural network," *IEEE Trans. Ind. Inform.*, vol. 19, no. 11, pp. 10 831–10 841, 2023.
- [34] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan *et al.*, "Argoverse: 3d tracking and forecasting with rich maps," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8748–8757.

- [35] S. Becker, R. Hug, W. Hübner, and M. Arens, “An evaluation of trajectory prediction approaches and notes on the trajnet benchmark,” *arXiv preprint arXiv:1805.07663*, 2018.
- [36] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [37] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, “Social gan: Socially acceptable trajectories with generative adversarial networks,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 2255–2264.
- [38] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” *Proc. Adv. Neu. Inf. Process. Syst.*, vol. 27, 2014.
- [39] Z. Chang, G. A. Koulouris, and H. P. H. Shum, “On the design fundamentals of diffusion models: A survey,” *arXiv*, 2023.
- [40] T. Gu, G. Chen, J. Li, C. Lin, Y. Rao, J. Zhou, and J. Lu, “Stochastic trajectory prediction via motion indeterminacy diffusion,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 17 113–17 122.
- [41] W. Mao, C. Xu, Q. Zhu, S. Chen, and Y. Wang, “Leapfrog diffusion model for stochastic trajectory prediction,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 5517–5526.
- [42] Y. Ma, X. Zhu, S. Zhang, R. Yang, W. Wang, and D. Manocha, “Trafficpredict: Trajectory prediction for heterogeneous traffic-agents,” in *Proc. AAAI Conf. Art. Intel.*, vol. 33, no. 01, 2019, pp. 6120–6127.
- [43] C. Yu, X. Ma, J. Ren, H. Zhao, and S. Yi, “Spatio-temporal graph transformer networks for pedestrian trajectory prediction,” in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 507–523.
- [44] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Proc. Adv. Neu. Inf. Process. Syst.*, vol. 30, 2017.
- [45] Y. Yuan, X. Weng, Y. Ou, and K. M. Kitani, “Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9813–9823.
- [46] C. Xu, M. Li, Z. Ni, Y. Zhang, and S. Chen, “Groupnet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning,” *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pp. 6488–6497, 2022.
- [47] P. Xu, J.-B. Hayet, and I. Karamouzas, “Socialvae: Human trajectory prediction using timewise latents,” in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 511–528.
- [48] C. Xu, W. Mao, W. Zhang, and S. Chen, “Remember intentions: Retrospective-memory-based trajectory prediction,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 6488–6497.
- [49] D. Li, Q. Zhang, S. Lu, Y. Pan, and D. Zhao, “Conditional goal-oriented trajectory prediction for interacting vehicles,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 12, pp. 18 758–18 770, 2024.
- [50] K. Sohn, H. Lee, and X. Yan, “Learning structured output representation using deep conditional generative models,” *Proc. Adv. Neu. Inf. Process. Syst.*, vol. 28, 2015.
- [51] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker, “Desire: Distant future prediction in dynamic scenes with interacting agents,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 336–345.
- [52] C. Wong, B. Xia, Z. Hong, Q. Peng, W. Yuan, Q. Cao, Y. Yang, and X. You, “View vertically: A hierarchical network for trajectory prediction via fourier spectrums,” in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 682–700.
- [53] J. Fang, C. Zhu, P. Zhang, H. Yu, and J. Xue, “Heterogeneous trajectory forecasting via risk and scene graph learning,” *IEEE Trans. Intell. Transp. Syst.*, 2023.
- [54] K. Guo, W. Liu, and J. Pan, “End-to-end trajectory distribution prediction based on occupancy grid maps,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 2242–2251.
- [55] H. Bi, Z. Fang, T. Mao, Z. Wang, and Z. Deng, “Joint prediction for kinematic trajectories in vehicle-pedestrian-mixed scenes,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 10 383–10 392.
- [56] T. Phan-Minh, E. C. Grigore, F. A. Boulton, O. Beijbom, and E. M. Wolff, “Covertnet: Multimodal behavior prediction using trajectory sets,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 14 074–14 083.
- [57] S. Chen and J. Wang, “Heterogeneous interaction modeling with reduced accumulated error for multiagent trajectory prediction,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 6, pp. 8040–8052, 2024.
- [58] J. Zhang, J. Yao, L. Yan, Y. Xu, and Z. Wang, “Sparse multi-relational graph convolutional network for multi-type object trajectory prediction,” in *Int. Joint Conf. Artif. Intell.*, 2024, pp. 1697–1705.
- [59] T. Fernando, S. Denman, S. Sridharan, and C. Fookes, “Soft+ hardwired attention: An lstm framework for human trajectory prediction and abnormal event detection,” *Neural networks*, vol. 108, pp. 466–478, 2018.
- [60] H. Xue, D. Q. Huynh, and M. Reynolds, “Poppl: Pedestrian trajectory prediction by lstm with automatic route class clustering,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 1, pp. 77–90, 2020.
- [61] I. A. Lawal, F. Poiesi, D. Anguita, and A. Cavallaro, “Support vector motion clustering,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 27, no. 11, pp. 2395–2408, 2016.
- [62] A. K.-F. Lui, Y.-H. Chan, and M.-F. Leung, “Modelling of destinations for data-driven pedestrian trajectory prediction in public buildings,” in *IEEE Int. Conf. Big Data*. IEEE, 2021, pp. 1709–1717.
- [63] I. Bae, J.-H. Park, and H.-G. Jeon, “Learning pedestrian group representations for multi-modal trajectory prediction,” in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 270–289.
- [64] Y. Wang, P. Zhang, L. Bai, and J. Xue, “Fend: A future enhanced distribution-aware contrastive learning framework for long-tail trajectory prediction,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 1400–1409.
- [65] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne,” *J. Mach. Learn. Res.*, vol. 9, no. 11, 2008.
- [66] N. Sai Madiraju, S. M. Sadat, D. Fisher, and H. Karimabadi, “Deep temporal clustering: Fully unsupervised learning of time-domain features,” *arXiv e-prints*, pp. arXiv:1802, 2018.
- [67] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [68] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *arXiv preprint arXiv:1412.3555*, 2014.
- [69] M. Cuturi and M. Blondel, “Soft-dtw: a differentiable loss function for time-series,” in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 894–903.
- [70] H. Sakoe and S. Chiba, “Dynamic programming algorithm optimization for spoken word recognition,” *IEEE transactions on acoustics, speech, and signal processing*, vol. 26, no. 1, pp. 43–49, 1978.
- [71] J. Zhao and L. Itti, “shapedtw: Shape dynamic time warping,” *Pattern Recognition*, vol. 74, pp. 171–184, 2018.
- [72] J. MacQueen *et al.*, “Some methods for classification and analysis of multivariate observations,” in *Proc. Fifth Berkeley Symp. Math. Statist. Probab.*, vol. 1, no. 14, 1967, pp. 281–297.
- [73] D. A. Reynolds *et al.*, “Gaussian mixture models,” *Encyclopedia of biometrics*, vol. 741, no. 659-663, 2009.
- [74] E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with gumbel-softmax,” *arXiv preprint arXiv:1611.01144*, 2016.
- [75] Y. Li, R. Liang, W. Wei, W. Wang, J. Zhou, and X. Li, “Temporal pyramid network with spatial-temporal attention for pedestrian trajectory prediction,” *IEEE Trans. Netw. Sci. Eng.*, vol. 9, no. 3, pp. 1006–1019, 2021.
- [76] J. Liang, L. Jiang, K. Murphy, T. Yu, and A. Hauptmann, “The garden of forking paths: Towards multi-future trajectory prediction,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 10 508–10 518.
- [77] Y. Zhang, W. Guo, J. Su, P. Lv, and M. Xu, “Bip-tree: Tree variant with behavioral intention perception for heterogeneous trajectory prediction,” *IEEE Trans. Intell. Transp. Syst.*, 2023.
- [78] Z. Zhou, L. Ye, J. Wang, K. Wu, and K. Lu, “Hivt: Hierarchical vector transformer for multi-agent motion prediction,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 8813–8823.
- [79] A. Monti, A. Bertugli, S. Calderara, and R. Cucchiara, “Dag-net: Double attentive graph neural network for trajectory forecasting,” in *Int. Conf. Pattern Recognit.* IEEE, 2021, pp. 2551–2558.
- [80] T. Zhao, Y. Xu, M. Monfort, W. Choi, C. Baker, Y. Zhao, Y. Wang, and Y. N. Wu, “Multi-agent tensor fusion for contextual trajectory prediction,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 12 126–12 134.
- [81] J. Liang, L. Jiang, and A. Hauptmann, “Simaug: Learning robust representations from simulation for trajectory prediction,” in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 275–292.
- [82] Y. Wang and S. Chen, “Multi-agent trajectory prediction with spatio-temporal sequence fusion,” *IEEE Trans. Multimedia*, 2021.
- [83] M. N. Azadani and A. Boukerche, “Capha: A novel context-aware behavior prediction system of heterogeneous agents for autonomous vehicles,” *IEEE Trans. Veh. Technol.*, 2023.
- [84] X. Chen, H. Zhang, Y. Hu, J. Liang, and H. Wang, “Vnagt: Variational non-autoregressive graph transformer network for multi-agent trajectory prediction,” *IEEE Trans. Veh. Technol.*, vol. 72, no. 10, pp. 12 540–12 552, 2023.

- [85] N. Rhinehart, K. M. Kitani, and P. Vernaza, “R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting,” in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 772–788.
- [86] S. H. Park, G. Lee, J. Seo, M. Bhat, M. Kang, J. Francis, A. Jadhav, P. P. Liang, and L.-P. Morency, “Diverse and admissible trajectory forecasting through multimodal context understanding,” in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 282–298.
- [87] C. Tang and R. R. Salakhutdinov, “Multiple futures prediction,” *Proc. Adv. Neu. Inf. Process. Syst.*, vol. 32, 2019.
- [88] X. Xu, W. Liu, and L. Yu, “Trajectory prediction for heterogeneous traffic-agents using knowledge correction data-driven model,” *Information Sciences*, vol. 608, pp. 375–391, 2022.
- [89] B. Yang, G. Yan, P. Wang, C.-Y. Chan, X. Song, and Y. Chen, “A novel graph-based trajectory predictor with pseudo-oracle,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 12, pp. 7064–7078, 2021.
- [90] I. Bae and H.-G. Jeon, “A set of control points conditioned pedestrian trajectory prediction,” in *Proc. AAAI Conf. Art. Intel.*, vol. 37, no. 5, 2023, pp. 6155–6165.
- [91] Y. Wu, L. Wang, S. Zhou, J. Duan, G. Hua, and W. Tang, “Multi-stream representation learning for pedestrian trajectory prediction,” in *Proc. AAAI Conf. Art. Intel.*, vol. 37, no. 3, 2023, pp. 2875–2882.
- [92] C. Xu, R. T. Tan, Y. Tan, S. Chen, Y. G. Wang, X. Wang, and Y. Wang, “Eqmotion: Equivariant multi-agent motion prediction with invariant interaction reasoning,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 1410–1420.
- [93] I. Bae, J. Oh, and H.-G. Jeon, “Eigentrajjectory: Low-rank descriptors for multi-modal trajectory forecasting,” *arXiv preprint arXiv:2307.09306*, 2023.
- [94] L. Shi, L. Wang, S. Zhou, and G. Hua, “Trajectory unified transformer for pedestrian trajectory prediction,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 9675–9684.
- [95] Y. Dong, L. Wang, S. Zhou, and G. Hua, “Sparse instance conditioned multimodal trajectory prediction,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 9763–9772.
- [96] B. Yang, F. Fan, R. Ni, H. Wang, A. Jafaripournimchahi, and H. Hu, “A multi-task learning network with a collision-aware graph transformer for traffic-agents trajectory prediction,” *IEEE Trans. Intell. Transp. Syst.*, 2024.
- [97] C. Xu, Y. Wei, B. Tang, S. Yin, Y. Zhang, S. Chen, and Y. Wang, “Dynamic-group-aware networks for multi-agent trajectory prediction with relational reasoning,” *Neural Networks*, vol. 170, pp. 564–577, 2024.
- [98] F. Marchetti, F. Becattini, L. Seidenari, and A. Del Bimbo, “Smemo: social memory for trajectory forecasting,” *IEEE Trans. Pattern Anal. Mach. Intell.*, 2024.
- [99] R. Liang, Y. Li, J. Zhou, and X. Li, “Stglow: A flow-based generative framework with dual-graphormer for pedestrian trajectory prediction,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 11, pp. 16 504–16 517, 2024.
- [100] Y. Peng, G. Zhang, J. Shi, X. Li, and L. Zheng, “Mrtraj: A novel non-autoregressive approach for human trajectory prediction,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 4, pp. 2318–2331, 2024.
- [101] S. Kim, H.-g. Chi, H. Lim, K. Ramani, J. Kim, and S. Kim, “Higher-order relational reasoning for pedestrian trajectory prediction,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 15 251–15 260.
- [102] X. Lin, T. Liang, J. Lai, and J.-F. Hu, “Progressive pretext task learning for human trajectory prediction,” in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 197–214.
- [103] J. Amirian, J.-B. Hayet, and J. Pettré, “Social ways: Learning multimodal distributions of pedestrian trajectories with gans,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2019, pp. 0–0.
- [104] W. Chen, H. Sang, J. Wang, and Z. Zhao, “Wtgcnn: wavelet transform graph convolution network for pedestrian trajectory prediction,” *Int. J. Mach. Learn. Cybern.*, pp. 1–18, 2024.
- [105] —, “Iggcn: Individual-guided graph convolution network for pedestrian trajectory prediction,” *Digital Signal Processing*, vol. 156, p. 104862, 2025.
- [106] J. Li, L. Yang, Y. Chen, and Y. Jin, “Mfan: Mixing feature attention network for trajectory prediction,” *Pattern Recognition*, vol. 146, p. 109997, 2024.
- [107] A. Gu, K. Goel, and C. Ré, “Efficiently modeling long sequences with structured state spaces,” *arXiv preprint arXiv:2111.00396*, 2021.
- [108] A. Gu and T. Dao, “Mamba: Linear-time sequence modeling with selective state spaces,” *arXiv preprint arXiv:2312.00752*, 2023.
- [109] M. Ronen, S. E. Finder, and O. Freifeld, “Deepdpm: Deep clustering with an unknown number of clusters,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 9861–9870.
- [110] A. Xiao, H. Chen, T. Guo, Q. Zhang, and Y. Wang, “Deep plug-and-play clustering with unknown number of clusters,” *Trans. Mach. Learn. Res.*, 2022.
- [111] P. Lv, W. Wang, Y. Wang, Y. Zhang, M. Xu, and C. Xu, “Ssagcn: Social soft attention graph convolution network for pedestrian trajectory prediction,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 9, pp. 11 989–12 003, 2024.



Ruochen Li is a PhD candidate in the Department of Computer Science at Durham University, UK. His research interests include trajectory prediction, computer vision, time-series analysis, and graph convolutional networks.



Stamos Katsigiannis (Member, IEEE) received the B.Sc. (Hons.) degree in informatics and telecommunications from the National and Kapodistrian University of Athens, Greece, in 2009, the M.Sc. degree in computer science from the Athens University of Economics and Business, Greece, in 2011, and the Ph.D. degree in computer science (biomedical image and general-purpose video processing) from the National and Kapodistrian University of Athens, Greece, in 2016. He is currently an Associate Professor with the Department of Computer Science, Durham University, UK. He has participated in multiple national and international research projects and has authored and co-authored over 70 research publications. His research interests include machine learning, natural language processing, affective computing, image analysis, and image and video quality.



Tae-Kyun (T-K) Kim is a Professor and the director of Computer Vision and Learning Lab at School of Computing, KAIST, since 2020, and has been a Reader at Imperial College London, UK in 2010–2024. He obtained his PhD from Univ. of Cambridge in 2008 and Junior Research Fellowship (governing body) of Sidney Sussex College, Univ. of Cambridge during 2007–2010. His BSc and MSc are from KAIST. His research interests lie in machine (deep) learning for 3D vision and generative AI, he has co-authored over 100 academic papers in top-tier conferences and journals in the field. He was the general chair of BMVC17 in London, and the program co-chair of BMVC23, Associate Editor of Pattern Recognition Journal, Image and Vision Computing Journal. He regularly serves as an Area Chair for the top-tier AI/vision conferences.



Hubert P. H. Shum (Senior Member, IEEE) is a Professor of Visual Computing and the Director of Research for the Department of Computer Science at Durham University, specialising in modelling spatio-temporal information with responsible AI. He is also a Co-Founder and the Co-Director of Durham University Space Research Centre. Before this, he was an Associate Professor at Northumbria University and a Postdoctoral Researcher at RIKEN Japan. He received his PhD degree from the University of Edinburgh. He chaired conferences such as Pacific Graphics, BMVC and SCA. He has authored over 200 research publications in the fields of Computer Vision, Computer Graphics and AI in Healthcare, underpinned by Responsible AI designs and algorithms.