

SIMMF: Semantics-aware Interactive Multiagent Motion Forecasting for Autonomous Vehicle Driving

Vidyaa Krishnan Nivash and Ahmed H. Qureshi

Abstract—Autonomous vehicles require motion forecasting of their surrounding multiagents (pedestrians and vehicles) to make optimal decisions for navigation. The existing methods focus on techniques to utilize the positions and velocities of these agents and fail to capture semantic information from the scene. Moreover, to mitigate the increase in computational complexity associated with the number of agents in the scene, some works leverage Euclidean distance to prune far-away agents. However, distance-based metric alone is insufficient to select relevant agents and accurately perform their predictions. To resolve these issues, we propose the Semantics-aware Interactive Multiagent Motion Forecasting (SIMMF) method to capture semantics along with spatial information and optimally select relevant agents for motion prediction. Specifically, we achieve this by implementing a semantic-aware selection of relevant agents from the scene and passing them through an attention mechanism to extract global encodings. These encodings along with agents’ local information, are passed through an encoder to obtain time-dependent latent variables for a motion policy predicting the future trajectories. Our results show that the proposed approach outperforms state-of-the-art baselines and provides more accurate and scene-consistent predictions. The demonstration video is available at: <https://youtu.be/Wjla071BPBA>

I. INTRODUCTION

Driving is a deceptively simple activity requiring continuous visual sampling, periodic control, and sustained attention [1]. Humans simplify this task by sampling the environment for critical control cues, typically involving predictable extrapolations of well-learned circumstances. Human cognition can perform these perceptual estimations with high accuracy to ensure safe driving. According to [2], two separate visual perceptual mechanisms exist to make good estimations. The first being the ambient mode, concentrates on factors such as spatial orientation and locomotion. It involves vehicle guidance using information gained from optical flow to derive estimations of vehicle positioning. The second being the focal mode, focuses on object recognition, identification, and the semantics associated with it. It concentrates on the detection and identification of objects of importance in the environment such as physical barriers, and roadway threats, and uses its practical semantic associations as critical cues for deriving estimations.

The current work in designing autonomous driving motion prediction methods focuses explicitly on the ambient mode and takes only implicit cues from the second, i.e., the focal mode, by inherently learning from data. Hence, the lack of context from the explicit modeling of the focal mode and

its inherent interaction with the ambient mode limits the existing methods’ performance accuracy and efficiency in complex environments. Therefore, in this paper, we present SIMMF, an extensible approach to introduce the focal mode and imbibe an active interaction between the two modes of perceptual mechanisms. Our key contributions are three-fold: Firstly, we formalize focal mode by formulating an algorithm to extract the semantics of agents and scene elements from environment maps. This extracted information is then utilized to optimize the grouping of relevant agents in a scene for their scene-consistent motion forecasting. Secondly, we combine both modes of perceptual systems through an attention mechanism to obtain their global, interactive representation. Thirdly, we capture inter-group agents’ interactions and new relevances in the data in a time-series manner using temporal encodings of focal and ambient representations. Finally, the motion prediction policy uses these temporal encodings to predict agents’ future trajectories under scene consistency. Hence, SIMMF aims to inculcate an active interaction between the two visual systems and induce an ability to complement one another, achieving mutual reinforcement for trajectory predictions. This results in an overall reduction in response time and improves the accuracy of trajectory predictions compared to state-of-the-art (SOTA) baseline methods.

II. RELATED WORK

Research in trajectory prediction has explored multiple facets of driving, uncovering a broad spectrum of approaches and possibilities. Early works majorly delved into decision-making processes based on deterministic models like Dynamic Bayesian Networks [3], Gaussian Process Regression (GPR) [4] and Recurrent Neural Networks (RNNs) [5] [6] [7]. According to the review papers [8] [9], motion prediction techniques can be classified based on the number of modes, level or type of abstraction, and method of approach. Many data-based approaches in the literature rely on LSTMs or simpler CNNs [10] to learn from agent states and histories. But these techniques do not ensure scene consistency, i.e., the trajectory predictions of multiple agents collide with each other or with static obstacles. To resolve this issue, recent studies have employed conditional variational autoencoders (CVAEs) or more recent methods such as generative adversarial networks (GANs) and attention mechanisms to achieve SOTA work in motion prediction for autonomous vehicles.

However, these techniques have time-independent latent variables and hence do not capture temporal dependencies within each agent’s history and dynamic relationships be-

The authors are affiliated with the Purdue University, West Lafayette, IN 47907, USA. Emails: {krish255, ahqureshi}@purdue.edu

tween agents. For example, in CVAE, the latent variable encodings for the input agent states and histories are constant with respect to time. Hence they do not explicitly learn the temporal graph governing the causal relationships in its inputs. The papers [11] [12] state this problem for general time-series data and proposes temporally dynamic latent variables to account for the lag in its memory states.

Some authors tend to resolve this issue using RNN for encoding the temporal information [13] [7], but these methods cannot handle multimodal trajectory predictions. Similarly, [14] predicts trajectories based on lane graphs by exploring both lateral (e.g.: lane keeping, turning) and longitudinal (e.g.: acceleration, braking) uncertainties. However, they do not capture the interaction between the two modes. Also, such works [15] focus on lane semantics, thus the other semantics (roadblocks, pedestrian walkways, barriers) are left unexplored.

Another recent and perhaps the most relevant method to our approach is ScePT [16] framework, which follows an ambient mode of perception, ensuring scene consistency and computational efficiency. It is a discrete CVAE model that outputs joint trajectory predictions and achieves scene consistency using an explicit collision penalty as regularization during training. Furthermore, it reduces the computational complexity by forming subgraphs of agents, also known as cliques, based on Euclidean distance and jointly predicts their future states rather than for an individual agent or a whole set of agents. Although ScePT provides SOTA performance, it lacks the following features that limit its performance. First, it does not account for the temporal dependencies of the agents, which, as highlighted earlier, is crucial for better performance. Second, it only considers the ambient mode of perception and does not have a focal mode based on agents’ local perception fields. Lastly, it forms cliques based on Euclidean distance only and does not account for other scene semantics such as lanes, road barriers, etc. Therefore, in this paper, we propose SIMMF, a full-stack motion prediction model that, with ambient and focal modes of perception, considers all possible scene semantics and temporal dependencies of selected agents, as summarized in Table I, and provides improved performance among all baselines.

Method	SC	TE	AF	SI	OC
SRNN[13]					✓
SSTran[12]	✓	✓			
LaFormer[15]		✓		✓	
LaneGraph[14]		✓		✓	
ScePT[16]	✓				✓
SIMMF	✓	✓	✓	✓	✓

TABLE I: Legend: SC - Scene Consistency, TE - Temporal Encoding, AF - Ambient and focal mode interaction, SI - Semantic-aware Interaction, OC - Optimal computation time

III. PROPOSED METHOD

Our work presents a method to generate joint trajectory predictions of all the interacting agents in a scene. These agents are captured as a spatiotemporal graph $G^t = (V, E)$

for each timestep t from the scene, where V represents the nodes and E represents the graph’s edges. The agents are represented as nodes, and the interactions between them are represented by the edges. Henceforth, the terms nodes and agents are used interchangeably and are broadly classified into two types, namely vehicles and pedestrians utilized throughout the model to determine the dynamics ψ of each agent based on its respective semantics. The state of a node is represented by $s \in \mathbb{R}^4$ containing information about the position $p = (x, y)$, orientation θ , and velocity (v_x, v_y) . Specifically, s_v contains vehicle state information $s_v = (x, y, \nu, \theta)$, where ν is the L_2 norm of (v_x, v_y) and s_p contains pedestrian information $s_p = (x, y, v_x, v_y)$. The states are stacked together from past H timesteps to timestep t , given by $\vec{h}_s^t = (s^{t-H}, s^{t+1-H}, \dots, s^t)$ while \vec{h}_e^t represents the corresponding stack of edge information. We use the state information to model an edge between a pair of nodes. The edge quantifies the degree of interaction ($\alpha_{ij} \in [0, 1]$), where (i, j) denotes the pair of nodes.

Our work assumes that maps containing geometric and semantic information are available. These maps are converted into local maps around each node. For a particular pixel resolution, we consider a corresponding fixed $K \times K$ size for all local maps, each modified into having L semantic channels given by $M^t \in \mathbb{R}^{K \times K \times L}$. The map semantic channels are constituted by the token identifiers τ . For example, roadblocks are defined by a set of edge line tokens $\tau_{e0:N}$ where N represents the total number of tokens required to define the spline. Similarly, lanes are represented by curve spline tokens τ_l . We form cliques C^t , i.e., subgraphs of agents having a high α_{ij} for a set of sampled timesteps $t \in (0, T)$, T representing the total number of timestamps considered for clique formation. We form the cliques based on the state and edge histories along with map semantics given by $C_{1:n}^t = \sigma(s_{t1:N}, \vec{h}_s^t, \vec{h}_e^t, M^t)$, where n represents the number of cliques formed from N number of agents in the scene. The function σ leverages various road semantics from M^t , at time t , such as lane barriers, crosswalks, lane directions, the distance between agents, etc., to quantify the edge value α between nodes and forms cliques of nodes with α greater than a threshold. Finally, given a clique, $C^{t:t-H}$, agents’ states, h_s^t , and interactions, h_e^t , at time t , we aim for a trajectory prediction function that predicts the future trajectories of each agent $s_{pred}^{t:t+\eta}$ for a horizon length η based on combined measures of the focal and ambient mode of visual perception.

A. Semantics-aware Clique Formation

The task of predicting the trajectories of all agents in a scene without collisions is complicated due to factors such as varying entry and exit timestamps and the interdependence between the multi-modal trajectories of different agents. To ensure scene consistency, each agent should be a part of the joint latent distribution, than individual predictions independent of its neighbors i.e. each agent should be associated with a discrete latent variable with cardinality N . Consequently, the cardinality N grows exponentially with the number of

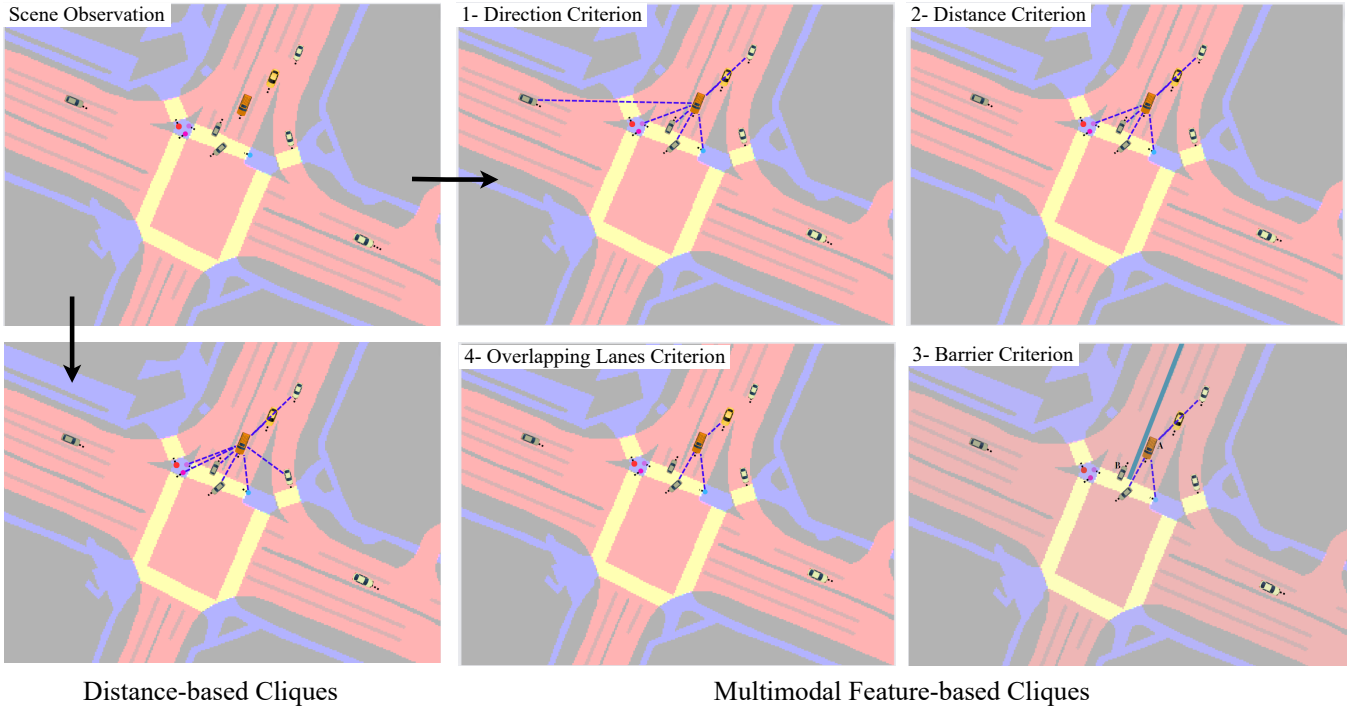


Fig. 1: Shows a step-by-step illustration of how the cliques are pruned using semantic information in order to capture the focal mode of visual perception. The orange vehicle represents the ego-vehicle (A), differently colored circles represent pedestrians, and the cars represent other vehicle agents. Fig 1(1) shows the direction criterion effectively pruning agents which move away from the ego-vehicle. Fig 1(2) shows how cliques are formed based on a distance threshold from A . Fig 1(3) represents the barrier criterion pruning agent B based on road semantics. Fig 1(4) illustrates how the lane semantics affect the clique formation.

nodes in the scene. To solve this increase in complexity, the scene can be broken down into smaller groups, called cliques formed based on degree of interaction. Fig 1 shows techniques to form cliques based on various criteria, such as the distance, direction, and semantics between the agents. These techniques prune the cliques to include only the most relevant agents for decision-making, thereby eliminating cliques that contain redundant information.

1) *Euclidean distance parameter*: Given a scene with multiple nodes, a spatio-temporal scene graph is generated where nodes represent agents and edges represent the interactions between these agents [16]. Each agent's closest future distance is ascertained as a measure of interaction, propagating forward each node i , according to an action a_i with a constant velocity model $\phi_{a_i}^t$. The $\phi_{a_i}^t$ maps the initial state of agent, i , to t time steps ahead into the future states based on a_i executed at the initial state, at time $t = 0$. The closest future distance between two agents is defined as $d_{ij} = \min \|\phi_{a_i}^t(s_i), \phi_{a_j}^t(s_j)\|_2$. We then define α_{ij} to populate the scene graph adjacency matrix as follows: the $\alpha_{ij} = 0$ when $d_{ij} > d_0$ and $\alpha_{ij} = \frac{d_0}{d_{ij}}$ when $d_{ij} < d_0$, where d_0 is the threshold distance fixed according to the agent type.

2) *Interaction detection using direction*: To ensure that the cliques do not include agents with minimal or no interaction, we categorize the agents based on whether they are moving toward or away from each other. For example, even if agents are within a distance threshold d_0 , they will

have no impact on each other if they are moving in opposite or perpendicular directions. One way to compute this is by taking the ego agent's viewpoint as the observer having position and velocity (p_o, v_o) and determining (p_i, v_i) of the other agents relative to the observer's frame of reference. The interaction between the ego vehicle and any other agent is considered relevant when there is a timestep within η such that the two velocity vectors intersect or lie within a relative distance threshold D [17]. Formally, $(p_o + tv_o) - (p_i + tv_i) = \vec{0} \iff V_o \cap V_i$. Since, $(p_o + tv_o) - (p_i + tv_i) = (p_o - p_i) + t(v_o - v_i)$, we check if the $\|(p_o - p_i) + t(v_o - v_i)\|_2 < D$. This technique also implicitly ensures scene consistency by introducing direction-based collision checking during clique formation.

3) *Barrier detection*: As mentioned earlier, the focal mode of visual cognition involves understanding the meaning and significance of the perceived elements in the environment as the first step for determining the next action. Specifically, humans imbibe critical cues from the perceived elements based on their knowledge, past experiences, and language understanding. For example, according to human cognition, it is trivial to conclude that two agents can not actively interact with each other if there are physical barriers present between them. But this intuitive fact is based on semantic information about the physical barriers. Hence, we implement this technique to detect barriers between agents to avoid clique formation between completely irrelevant agents. To understand the formation of cliques better, let us consider

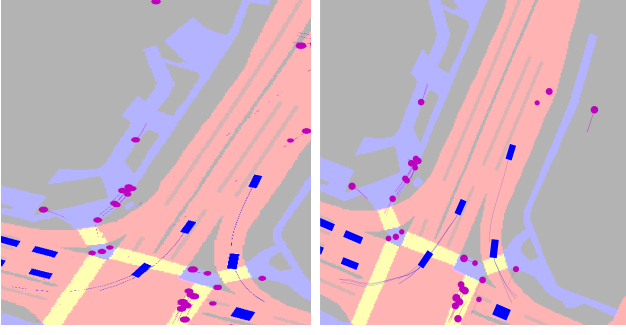


Fig. 2: Comparison of ground truth (left) and prediction trajectories (right) for a particular scene for both vehicles (blue rectangle) and pedestrians (pink circle). We can observe that each agent has multimodal predictions, with high-confidence predictions aligning with the ground truth paths.

the scene in Figure 1(3). In this scene, agents A and B are moving toward each other, but there is no possible interaction between the agents due to the physical barrier (road divider) between them. Hence these agents need not be considered together in the same clique. To check if there are any barriers between the two agents, map semantics can be used which are derived from scene segmentation. Given these map attributes as line segments dividing roads, we check if there are dividers between any two agents. Consider a line segment $\tau_{A \rightarrow B}$ drawn between the agents A and B, $\tau_{C \rightarrow D}$ drawn between C and D in Figure 1(3) respectively. The line segments $\tau_{A \rightarrow B}$ and $\tau_{C \rightarrow D}$ intersect if ACD and BCD have opposite directions meaning either ACD or BCD is counterclockwise but not both. To check the direction of ACD , we calculate the slopes m_{AC} and m_{AD} of AC and AD , respectively. If $m_{AC} < m_{AD}$, then ACD has a clockwise direction, else otherwise. We use this condition to check if two agents will intersect with each other.

4) *Overlapping lanes*: One of the techniques involved in the focal mode of visual cognition includes estimating the future possible trajectories based on lane semantics. In other words, the interaction between agents can also be determined based on the lanes which are available in the future for each agent. Based on this information, we can track if there will be any possible intersection of lanes between two agents for them to have an interaction. Lanes for each agent can be found using lane tracking algorithms [18] [19]. They can also be derived from map information assuming the data is accurate. Considering the map semantics to be accurate, we track the possible lanes of each agent. We have lane dividers spanning all the lanes on the roads. These dividers are defined by a set of token identifiers $\tau_{e_{0:N}}$ and are connected with each other by common end tokens τ_{e_0} and τ_{e_N} . The direction of traffic in a particular lane divider can be obtained by extrapolating any of the tokens from $\tau_{e_{0:N}}$. Each agent can be associated with a lane divider τ_{d_i} based on its current position. To find if any two agents would meet in the future, we find if there's any intersection in the list of possible lanes of each agent τ_{l_i} . In order to find τ_{l_i} for each agent i , we collect all lane dividers connected to τ_{d_i} by considering the dividers having τ_{e_N} as a common token. Consequently, we




Fig. 3: Graph is generated with nodes as agents. Cliques are formed based on various techniques mentioned in Section 3.1, further passed on to get batch cliques containing node histories and local maps. Upon further processing, they are fed into the Multimodal attention mechanism to get global and local encodings. These encodings are fed into a policy network to obtain closed-loop trajectory predictions.

check for the intersection of the list of lane tokens between agents, $\alpha_{ij} = \mathbb{I}(\tau_{l_i} \cap \tau_{l_j})$.

B. Multi-head Attention-based Global Encoding

Once the cliques are formed, their local semantic maps and node state and edge history encodings are extracted and further processed through a multi-head attention mechanism to obtain global encodings. To understand the significance of this mechanism, let us consider the state-of-the-art architecture for trajectory prediction consisting of a CVAE model that encodes node history and local map information through a Gibbs distribution and passes through a policy network to generate trajectory predictions. Since agents follow a regular pattern and are highly correlated with each other, global map and state history information is crucial to capture the joint trajectory predictions. Specifically, in cases where agents are occluded from the ego vehicle (cross-traffic scenario), local map information and its associated state history information are not sufficient to capture the global information of the agents in the clique. Also, global context is required to capture the inter-clique information essential for minimizing the collision rate between cliques. Hence, we propose a multi-head attention mechanism by passing map encodings, node history, and edge history information as unique heads through a custom-designed transformer network. The model captures different dependencies and relationships between these heads, thereby enabling active interaction between ambient and focal modes of the vision system. Let the map M_i^t of each agent i be encoded via CNN to produce local encodings $M_{l_i}^t$, and similarly, state and edge histories $(h_{s_i}^t, h_{e_i}^t)$ be encoded into $(h_{s_{l_i}}^t, h_{e_{l_i}}^t)$ using LSTMs. The multi-head attention mechanism acts on these local features $(M_{l_i}^t, h_{s_{l_i}}^t, h_{e_{l_i}}^t)$ to produce global features $(M_{g_i}^t, h_{s_{g_i}}^t, h_{e_{g_i}}^t)$ and appended with the local contextual information to form $(\bar{M}_{l_i}^t, \bar{h}_{s_{l_i}}^t, \bar{h}_{e_{l_i}}^t)$. The multi-head attention mechanism structure consists of a fully connected layer FC outputting the weights of an attention matrix A populated on its non-diagonal elements. In greater detail, FC layer outputs a probability weight for all agents with respect to each and every other agent. The local features are then multiplied with the weights to get a global feature vector of the same dimension, $M_{g_i}^t = \sum w_{i,j} M_{l_i}^t$, $h_{s_{g_i}}^t = \sum w_{i,j} h_{s_{l_i}}^t$ and $h_{e_{g_i}}^t = \sum w_{i,j} h_{e_{l_i}}^t$.

C. Temporal Latent Encodings

In SIMMF, we present a method that utilizes both modes of interaction to capture dynamic information by associating each node with a time-dependent latent variable. To achieve this, we encode the global information along with local information from both modes of visual perception into discrete latent variables $z_i^t = [z_1^t, z_2^t, z_3^t, \dots, z_N^t]$ for each clique $C_{1:n}^t$ consisting of agents 1 to N . To elaborate, consider the temporal encoding framework as a function g . We can write $z_i^t = g(\vec{h}_{s_i}^t, \vec{h}_{e_i}^t, \vec{M}_i^t, s_{ref_i}^{t:t+\eta})$ where s_{ref} denotes the reference trajectories during training, and the former variables represent the local and global encodings concatenated together. Note that this is contrary to recent work that captures temporal information implicitly by incorporating influence from neighboring nodes using edge encoders [20], which incorporates only the ambient mode. Consequently, the temporal information is ignored if the latent encodings are not iteratively captured based on timestamps, hurting prediction accuracy. The addition of the time-dependent latent variables has significantly improved the optimality in the choice of agents, especially in dense scenes.

D. Overall pipeline

In SIMMF, we generate a spatio-temporal graph G^t , with nodes representing agent and edges representing their interactions evaluated using the techniques mentioned in Section 3.1. The degree of interaction α is populated in an adjacency matrix to form optimal cliques C^t , i.e., sub-graphs of agents which have considerable interaction with each other. These cliques are further processed to obtain the active nodes and their relative states in each clique. Consequently, they are rearranged into a list of active nodes containing their corresponding relative state and edge histories, i.e., the batches are populated with the aggregated information of all the nodes along with their corresponding clique details. Further, we generate local semantic maps [21] M^t , by rotating and cropping according to each agent’s heading and type. These local maps are passed on to CNN to get local map encodings M_l^t , while the state and edge histories are encoded into the feature vectors $(h_{s_i}^t, h_{e_i}^t)$ via LSTMs. These vectors are passed through the Multi-head attention-based mechanism to obtain global encodings $(M_g^t, h_{s_g}^t, h_{e_g}^t)$. The global encodings appended with the local encodings to form $(\vec{M}^t, \vec{h}_s^t, \vec{h}_e^t)$ are further passed through the Gibbs discrete sampling to produce latent variable encodings for each timestamp $0 : T$ to capture the temporal information. These time-dependent latent variables z^t are passed through the decoder containing the policy network to obtain the joint trajectory predictions $s_{pred}^{t:t+\eta}$ (Figure 3).

IV. EXPERIMENTAL RESULTS

Our experiments use nuScenes [21] and Argoverse2 [22] datasets for evaluation.

A. Evaluation Metrics and Baselines

We use the standard evaluation metrics, also employed by the prior methods [15], [16], [20], [23] namely Average Displacement Error (ADE), Final Displacement Error (FDE),

and Miss Rate. Miss Rate denotes the number of scenarios where none of the predicted trajectories are within 2.0 meters of ground truth calculated at the final prediction. [22] We choose the most likely mode for our evaluations on both datasets, denoted as $minADE_1$ and $minFDE_1$ in Argoverse metrics [22]. The former is the average L2 distance between the ground truth and predicted trajectories, whereas the latter is the L2 distance between the best-predicted trajectory and its ground truth. In addition, we also introduce a new metric, named Mean Average Count (mAC) which is inspired by a metric in [24]. The mAC is determined as follows: We compute the trajectory prediction of the highest confidence by choosing the Best-of-N extension from the encoder output, i.e. we take the best from the N highest probability modes of the latent variable encodings taken at a time. Next, we check if the prediction is within lateral and longitudinal thresholds $\eta/3\mu(v_x)$ and $\eta/3\mu(v_y)$, respectively, where η is time length and μ is a scaling factor computed based on the velocities v_x and v_y . Finally, if the prediction lies within the thresholds, then it is considered a true positive; otherwise considered a false positive. The average of false positives is calculated per clique to estimate the accuracy of each clique. The overall mAC is calculated as the mean of all the cliques in the scene. For all metrics, the lower is better, and bold/italic marks the best/second-best value.

Method	@1s	@2s	@3s	@4s
S-LSTM [24]	0.47	-	1.61	-
CSP[25]	0.46	-	1.50	-
CAR-Net [26]	0.38	-	1.35	-
SpAGNN [27]	0.35	-	1.23	-
Trajectron++ [19]	0.07	0.45	<i>1.14</i>	2.20
ScePT(Best-of-3)	0.40	0.80	1.36	<i>2.14</i>
SIMMF(Best-of-3)	<i>0.32</i>	<i>0.48</i>	0.72	1.18

TABLE II: Comparison of FDE (in meters) with Baselines on nuScenes. Ours perform better among most of the baselines, specifically in the latter timesteps of the prediction horizon.

# samples	ADE		mAC	
	ScePT	Ours	ScePT	Ours
2	0.35/0.11	0.22/0.07	0.33	0
3	0.31/0.11	0.26/0.07	1.33	0
5	0.31/0.07	0.17/0.06	0.33	0.67
10	0.35/0.07	0.27/0.03	2	1.67

TABLE III: ADE (in meters), using the most likely mode and mAC (in percentage) on nuScenes. SIMMF performs better for all samples for ADE, for most samples for mAC.

Our results are compared against these baselines: Social LSTM [25] where each agent is modeled with LSTMs and followed by the multiagent ‘‘Social’’ pooling layer; Convolutional Social Pooling (CSP) [26] which explicitly considers a fixed number of trajectories for a classification task; CAR-Net [27] solves the prediction task by localizing the regions of interest for each agent using LSTMs-based visual attention; SpAGNN [28] where raw LIDAR data and semantic maps are encoded using CNN followed by GNN to produce probabilistic trajectories; Trajectron++ [20] and ScePT [16]

time	Ours	S	T	A	B	O	D	ScePT
0.5s	0.19	0.28	0.22	0.33	0.19	0.20	0.31	0.33
1.5s	0.41	0.50	0.46	0.55	0.43	0.44	0.50	0.54
2.5s	0.56	0.69	0.62	0.76	0.60	0.61	0.68	0.76
3.5s	0.93	1.08	0.99	1.21	0.92	0.92	0.76	1.22

TABLE IV: Ablation study of techniques presented in SIMMF based on FDE on nuScenes. SIMMF without the following methods indicated by legend: S- Semantic Cliques, T - Temporal encodings, A - Attn mechanism, B - Barrier Detection, O - Overlapping lanes, D - Interaction using direction. *A* has the highest influence on SIMMF, followed by the parameter *D* in semantic-aware clique formation.

where the agents are represented as a spatiotemporal graph and the predictions are forecasted through CVAE model.

B. Results

In this section, we present the comparison and ablation analysis on nuScenes [21] and argoverse [22] datasets. The ADE and FDE are represented in meters, and mAC is represented as the percentage of misses, i.e., the percentage of the average number of nodes having predictions that do not lie within the expected threshold region. Furthermore, we evaluate with Best-of-N metric similar to the baselines to establish a fair comparison.

Comparison with baselines: In this section, we compare

# samples	minADE		Miss Rate	
	ScePT	Ours	ScePT	Ours
2	1.58/1.12	1.47/0.96	0.66	0.58
3	1.47/1.01	1.45/0.98	0.64	0.59
5	1.46/ 0.96	1.37/0.97	0.63	0.55
10	1.46/1.01	1.38/0.98	0.65	0.58

TABLE V: minADE (in meters), using the most likely mode and Miss Rate on Argoverse. SIMMF performs better on most samples for ADE, all samples for Miss Rate.

our approach, SIMMF, against the given baselines. Table II summarizes the FDE-based comparison of all methods on nuScenes. SIMMF gives overall better results among almost all the baselines, especially in the latter timesteps of the prediction horizon. This clearly shows that the temporal dependencies in lateral encodings are vital in reducing the accumulated error from the previous timesteps' predictions. Although Trajectron++ [20] gives the best result in the short term, it is computationally expensive to execute as it does not form cliques. Among baselines, only the ScePT forms cliques; therefore, we center our remaining comparison analysis against the ScePT method. Table III compares our method with ScePT using mAC and ADE on nuScenes. The results of ADE (represented for vehicles/pedestrians respectively) show that our model performs better than ScePT. Furthermore, as we increase the number of samples taken from the Gibbs distribution, there is an increase in miss rate, possibly due to the impact of the tunable risk measure incorporated from ScePT that modifies the confidence level of predictions. Specifically, there are high-confidence predictions that are different from the ground truth when




Fig. 4: The incorporation of an attention mechanism in SIMMF demonstrates a reduced collision rate compared to ScePT on nuScenes, resulting in improved scene consistency.

increasing the number of sample encodings taken at a time. Table V compares our method with ScePT on Argoverse2. The results of minADE and Miss Rate show that SIMMF performs slightly better, with significant improvement seen at 2 samples. Particularly, SIMMF performs best for vehicles at 5 samples. In terms of collision rate, SIMMF shows significantly enhanced scene consistency than ScePT on nuScenes, as seen from the graphs in Figure 4, demonstrating a considerable decrease in collision rate, specifically with the use of the attention mechanism. The qualitative comparison of trajectory predictions against ground truth is visualized in Figure 2. It can be seen that our predictions closely align with the ground truth for both vehicles and pedestrians.

Ablation study: To develop an understanding of the individual components of SIMMF, we conduct an ablation study having two parts. First, we ablate the model architecture on semantic aware clique formation, temporal encodings, and attention mechanism. Secondly, we investigate the impact of semantic cliques by ablating the various techniques of Section 3.1, including direction, barrier detection, and overlapping lanes. We utilize ScePT for the ablation study since it quantifies the degree of interaction solely using the Euclidean distance parameter. Table IV summarizes our results, columns 2-5 and 5-9 highlight the first and second parts of our ablation study. It can be observed that the attention mechanism has the greatest impact on the accuracy, whereas the temporal factor has more impact on the later timesteps. From the second part, we can ascertain that the direction factor has the greatest impact.

V. LIMITATIONS AND FUTURE WORK

One of the limiting factors of our approach is that we assume road semantics information is available and accurate, which might not be the case in real-world settings. Therefore, our future work includes working on an end-to-end model using only camera inputs to provide solutions to vehicles that lack expensive onboard sensors for feature extraction. Secondly, our method implicitly captures inter-clique interaction, which with explicit modeling, may lead to further improvements in our performance. Therefore, another of our future works is explicitly capturing inter-clique interactions that may be achieved using reparameterization techniques to backpropagate through the time-dependent latent variables sampled during training for temporal latent encodings.

REFERENCES

- [1] C. Castro, *Human factors of visual and cognitive performance in driving*. CRC Press, 2008.
- [2] R. Tyrrell, K. Rudolph, B. Eggers, and H. Leibowitz, "Evidence for the persistence of visual guidance information," *Perception & psychophysics*, vol. 54, no. 4, pp. 431–438, 1993.
- [3] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *ROBOMECH journal*, vol. 1, no. 1, pp. 1–14, 2014.
- [4] J. M. Wang, D. J. Fleet, and A. Hertzmann, "Gaussian process dynamical models for human motion," *IEEE transactions on pattern analysis and machine intelligence*, vol. 30, no. 2, pp. 283–298, 2007.
- [5] J. Morton, T. A. Wheeler, and M. J. Kochenderfer, "Analysis of recurrent neural networks for probabilistic modeling of driver behavior," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 5, pp. 1289–1298, 2016.
- [6] A. Vemula, K. Muelling, and J. Oh, "Social attention: Modeling attention in human crowds," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4601–4607, IEEE, 2018.
- [7] X. Jia, L. Chen, P. Wu, J. Zeng, J. Yan, H. Li, and Y. Qiao, "Towards capturing the temporal dynamics for trajectory prediction: a coarse-to-fine approach," in *Conference on Robot Learning*, pp. 910–920, PMLR, 2023.
- [8] F. Leon and M. Gavrilescu, "A review of tracking and trajectory prediction methods for autonomous driving," *Mathematics*, vol. 9, no. 6, p. 660, 2021.
- [9] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras, "Human motion trajectory prediction: A survey," *The International Journal of Robotics Research*, vol. 39, no. 8, pp. 895–935, 2020.
- [10] N. Nikhil and B. Tran Morris, "Convolutional neural network for trajectory prediction," in *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pp. 0–0, 2018.
- [11] H. Hosseini, S. Kannan, B. Zhang, and R. Poovendran, "Learning temporal dependence from time-series data with latent variables," in *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 253–262, IEEE, 2016.
- [12] R. Girgis, F. Golemo, F. Codevilla, M. Weiss, J. A. D'Souza, S. E. Kahou, F. Heide, and C. Pal, "Latent variable sequential set transformers for joint multi-agent motion prediction," *arXiv preprint arXiv:2104.00563*, 2021.
- [13] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena, "Structural-rnn: Deep learning on spatio-temporal graphs," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5308–5317, 2016.
- [14] N. Deo, E. Wolff, and O. Beijbom, "Multimodal trajectory prediction conditioned on lane-graph traversals," in *Conference on Robot Learning*, pp. 203–212, PMLR, 2022.
- [15] M. Liu, H. Cheng, L. Chen, H. Broszio, J. Li, R. Zhao, M. Sester, and M. Y. Yang, "Laformer: Trajectory prediction for autonomous driving with lane-aware scene constraints," *arXiv preprint arXiv:2302.13933*, 2023.
- [16] Y. Chen, B. Ivanovic, and M. Pavone, "Scept: Scene-consistent, policy-based trajectory predictions for planning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17103–17112, 2022.
- [17] H. Herencia-Zapana, J.-B. Jeannin, and C. A. Munoz, "Formal verification of safety buffers for sate-based conflict detection and resolution," in *27th International Congress of the Aeronautical Sciences (ICAS 2010)*, no. NF1676L-9213, 2010.
- [18] V. Nguyen, H. Kim, S. Jun, and K. Boo, "A study on real-time detection method of lane and vehicle for lane change assistant system using vision system on highway," *Engineering science and technology, an international journal*, vol. 21, no. 5, pp. 822–833, 2018.
- [19] S. Sivaraman and M. M. Trivedi, "Integrated lane and vehicle detection, localization, and tracking: A synergistic approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 2, pp. 906–917, 2013.
- [20] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pp. 683–700, Springer, 2020.
- [21] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11621–11631, 2020.
- [22] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, *et al.*, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," *arXiv preprint arXiv:2301.00493*, 2023.
- [23] Y. Yuan, X. Weng, Y. Ou, and K. M. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9813–9823, 2021.
- [24] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou, *et al.*, "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9710–9719, 2021.
- [25] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 961–971, 2016.
- [26] N. Deo and M. M. Trivedi, "Multi-modal trajectory prediction of surrounding vehicles with maneuver based lstms," in *2018 IEEE intelligent vehicles symposium (IV)*, pp. 1179–1184, IEEE, 2018.
- [27] A. Sadeghian, F. Legros, M. Voisin, R. Vesel, A. Alahi, and S. Savarese, "Car-net: Clairvoyant attentive recurrent network," in *Proceedings of the European conference on computer vision (ECCV)*, pp. 151–167, 2018.
- [28] S. Casas, C. Gulino, R. Liao, and R. Urtasun, "Spatially-aware graph neural networks for relational behavior forecasting from sensor data," *arXiv preprint arXiv:1910.08233*, 2019.