Polar Collision Grids: Effective Interaction Modelling for Pedestrian Trajectory Prediction in Shared Space Using Collision Checks*

Mahsa Golchoubian[1†], Moojan Ghafurian[2], Kerstin Dautenhahn[2], Nasser Lashgarian Azad[1]

Abstract—Predicting pedestrians' trajectories is a crucial capability for autonomous vehicles' safe navigation, especially in spaces shared with pedestrians. Pedestrian motion in shared spaces is influenced by both the presence of vehicles and other pedestrians. Therefore, effectively modelling both pedestrian-pedestrian and pedestrian-vehicle interactions can increase the accuracy of the pedestrian trajectory prediction models. Despite the huge literature on ways to encode the effect of interacting agents on a pedestrian's predicted trajectory using deep-learning models, limited effort has been put into the effective selection of interacting agents. In the majority of cases, the interaction features used are mainly based on relative distances while paying less attention to the effect of the velocity and approaching direction in the interaction formulation. In this paper, we propose a heuristic-based process of selecting the interacting agents based on collision risk calculation. Focusing on interactions of potentially colliding agents with a target pedestrian, we propose the use of time-to-collision and the approach direction angle of two agents for encoding the interaction effect. This is done by introducing a novel polar collision grid map. Our results have shown predicted trajectories closer to the ground truth compared to existing methods (used as a baseline) on the HBS dataset.

I. Introduction

Autonomous vehicles (AVs) are envisioned to be part of our everyday future transportation system. For an efficient and safe operation of these AVs in spaces where pedestrians are present, they should be able to predict the future trajectories of nearby pedestrians.

Accurate pedestrian trajectory prediction requires the integration of different sources of information. This includes the past trajectory history of each agent as well as the between-agent interaction effects (e.g., how one pedestrian's movement may affect another pedestrian's). Due to the complexity of pedestrians' motion in interactive settings, many deep-learning models have been proposed in the literature for predicting pedestrian trajectories while considering the interaction effects (e.g.,[1], [2], [3], [4], [5], [6]). Most of these methods have focused their training on datasets that only include pedestrians without any vehicle present [1], [2], [3]. Thus, these past models have only accounted for pedestrian-pedestrian interaction effects. However, such

models cannot directly be used for predicting a pedestrian's trajectory in the presence of pedestrians and vehicles.

Recent work has started considering the effect of pedestrian-vehicle interaction in the prediction models as a distinct influencing factor [7], [8], [9], [10], [11], [12], [4]. Examples are using two separate pooling modules for information aggregation of pedestrians and vehicles that are interacting with the target pedestrian [7], or are accounting for different edge types in a graph neural network (GNN) architecture for distinguishing the two types of interaction apart [11], [4]. A complete review of these pedestrian trajectory prediction models that also account for pedestrian-vehicle interaction effects is provided in [13].

In deep learning approaches for modelling interaction effects, two main questions need to be answered: (1) How interacting agents are specified, and (2) How to formulate the effect of these interacting agents on the pedestrians' future trajectories. Most of the existing literature in this area has focused only on proposing different methods for answering the second question. Simplified answers have been given to the first question such as considering agents as interacting agents that are closer to each other than a specified threshold [1], [14], [6], [5], [15], [12], [16], or considering all the agents in the scene and aggregating all their trajectory information for formulating interaction [8], [7], [10]. While some proposed methods have considered learning the importance of each neighbouring agent through weightings learned in attention layers [8], [17], still, a measure for selecting the interacting agents in the first place is missing in the literature. For example, in classical heuristic-based methods, pedestrians' motions when encountering vehicles are designed to be governed by time-dependent collision risk features such as time to collision [18], [19] or minimum gap acceptance [20], [21] (defined as the time left between the pedestrian and the nearest upcoming vehicle).

Getting inspired by these methods, in this paper, we are proposing the use of Time To Collision (TTC) for specifying interacting agents (both pedestrians and vehicles) that can influence a target pedestrian's future trajectory.

Moreover, in response to the second question about formulating the possible effect of the interacting agents on a target pedestrian's trajectory, most of the existing models rely on relative distance features such as pooling the information of all interacting agents in a fixed grid map around the target pedestrian [1], [14], [3], [6], [5], or aggregating the information of the neighbouring agents using their relative distance to the target pedestrian [7], [2], [22], [17], [8]. In all these distance-dependent methods of capturing interaction,

the effects of approach directions and speed are missing while these two factors can indeed have a high impact on pedestrians' motion behaviour [23], [24]. In other words, two agents could be close, in terms of relative distance, but are not necessarily influencing each other's trajectory, if walking in parallel or getting away from each other.

Few other models [17], [12], [11] have considered the raw velocity information as an additional agent state. However, velocity information can be used for deriving more explainable engineered interaction features. Therefore, in this paper, we will take advantage of the information embedded in the velocity vectors to calculate the collision probabilities, and define a polar collision grid map for capturing interaction features that can influence a pedestrian's future trajectory.

We believe that a neural network model can learn interaction-influenced motion patterns more effectively if given more explainable engineered features as input.

We can summarize our contribution as follows:

1) Specifying interacting agents based on TTC as a more effective indicator compared to the relative distance used in the literature for studying interaction effect in pedestrian trajectory prediction
2) Introducing a novel polar collision grid map for capturing interacting effects in a deep-learning-based prediction model using both TTC and approaching angle information.

II. RELATED WORK

In deep learning-based pedestrian trajectory prediction models, an interaction encoder is responsible for building a feature that captures the between-agent interaction effects on the future trajectory. Various structures have been proposed for this interaction encoder in the literature which is reviewed here from two main perspectives: 1) The selection of interacting agents, 2) The encoded interaction features.

A. Specifying Interacting Agents

The deep learning methods used in the literature often rely on positional distance for reasoning about the interacting agents. In [1], [14], [6], [5] an occupancy grid map centred at each agent's position is used for constructing the interaction feature of that agent. With a fixed grid size, only agents closer than a threshold to that target agent are considered as having an interaction and are therefore considered for constructing that target agent's interaction feature, thus, only focusing on the effects caused by close-by agents.

In the GNN models proposed in [15], [12], [16], an agent's node is connected only to other nodes within a certain distance and the interaction effects of only those connected nodes are considered in the message passing process.

To overcome the limitation of these methods in capturing possible farther-away interactions, others have included the information of all agents in the scene for building the interaction feature [8], [7], or considered a fully connected graph in the GNN models used for trajectory prediction [10].

In a novel approach proposed by Li et al. (2021), the existence of an interaction edge in a GNN model is decided through a reinforcement learning (RL) framework on top of the prediction model [25]. In their RL formulation, the action is specified as making the edge between the two nodes on or off and the reward is defined based on how close the prediction of the selected graph connection is to the ground truth trajectory.

Su et al. (2022), have proposed different connectivity graphs for interaction consideration in their GNN structure namely *View graph*, *Direction graph* and *Rate graph* [26]. In the *View graph* only agents present in a target pedestrian's field of view are connected to that pedestrian's node in the graph and are considered to have an interaction effect on the pedestrians. For *Direction graph* and *Rate graph*, these edges are decided based on whether the two agents' moving direction has a crossover [26]. While the *Direction graph* proposed in [26] can potentially account for the collision risks, it should be pointed out that not every path crossover necessarily yields a collision. The agents on two crossing paths will not collide if the two agents get to that crossover point at different times. Therefore, we propose the direct use of collision checks for specifying the interacting agents while also accounting for the time alignments in the trajectories for defining a collision.

B. Interaction Features

Regardless of how the potentially interacting agents are selected, the interaction features used throughout the literature are again mainly constructed based on the relative distances of the neighbours to the target agent.

In [1], [14], [3], [6], [5], fixed-size grid maps centred at the target agents' locations is defined and an interaction tensor is built using the occupied cell of each neighbour according to relative distances to the target agent. Both rectangular shape grids [1], [14], [3] and circular shape (polar) grids [6], [5], [3] have been proposed. But even in the polar grids proposed in [6], [5], [3], only the angle of the distance vector between the two agents is accounted for in the grid construction without caring about the approach direction of the two agents.

In other models the relative distances between the target agent and its neighbours are embedded and concatenated with the hidden state of the neighbours and finally, this concatenated feature from all the neighbours is aggregated using a softmax [8] or a max pooling layer [7], [2], [22] to construct the interaction feature. In the soft attention mechanism used in [17] for aggregating the neighbours' hidden states, the bearing angle between the two agents is also used along with their distance in the calculation of the attention score for each neighbour.

Cheng et al. (2018) implemented a novel approach of using collision probability between each target agent and all its neighbours for building the interaction feature tensor inputted to an LSTM model [9]. But again their collision probability is calculated as an exponential function of the relative distance, ignoring the influence of speed.

Therefore, in our work, we are proposing a collision probability grid that stores actual TTC, and approaching angle information by using both velocity and position data.

III. METHOD

A. Problem Definition

Pedestrian trajectory prediction is defined as predicting the future positions of pedestrians given a short observed history of their own and the close-by agents' (pedestrians and vehicles) trajectories [1]. The position of pedestrians and vehicles at time step t along their trajectories are denoted as $X_t^{ped} = [X_t^1, X_t^2, ..., X_t^{n_p}]$ and $X_t^{veh} = [X_t^1, X_t^2, ..., X_t^{n_v}]$, where "ped" stands for pedestrian and "veh" stands for vehicle, and n_p and n_v are the numbers of pedestrians and vehicles in the scene respectively. The x-y position coordinate of pedestrian or vehicle i at time t in this formulation is defined as $X_t^i = (x_t^i, y_t^i)$. The velocity vector of each agent denoted as $V_t^i = (v_{x_t^i}, v_{y_t^i})$ can also be calculated using the position difference between successive time steps, knowing the timestamp between the captured frames.

Given the observed trajectories of all agents for the time steps between $0 < t \leq T_{obs}$, we predict the trajectory of the pedestrians across the future time period of $T_{obs} < t \leq T_{pred}$ as $\tilde{Y}_t = [\tilde{Y}_t^1, \tilde{Y}_t^2, ..., \tilde{Y}_t^{n_p}]$. The ground truth trajectory of pedestrians for any time step in this prediction period is denoted as $Y_t = [Y_t^1, Y_t^2, ..., Y_t^{n_p}]$. The predicted position coordinates $\tilde{Y}_t^i$ are derived from the predicted Gaussian distribution output of the model $\hat{Y}_t^i$, which will be described later in Eq. 6.

B. Overall Framework of the Trajectory Prediction Model

Our overall approach for predicting pedestrians' trajectories in an environment shared with both pedestrians and vehicles is shown in Fig. 1. The movement pattern in the pedestrian's trajectories is learned through an LSTM Network. The network receives three types of input features for each pedestrian: (1) the spatial embedded feature of the pedestrian, (2) the embedded pedestrian-pedestrian interaction feature, and (3) the embedded pedestrian-vehicle interaction feature. These features are aggregated together and used as the input to the LSTM at each timestep.

Similar to other works [1], [8], instead of predicting the actual future positions, our model will predict the parameters of a bi-variant Gaussian distribution of possible future positions to account for prediction uncertainty. Therefore, the output of our LSTM goes through a linear layer to predict the means (μ_x, μ_y), standard deviations (σ_x, σ_y) and the correlation coefficient (ρ) of the bi-variant Gaussian distribution.

1) Spatial embedded feature: The effect of each pedestrian's own trajectory history on the predicted next position is captured through this feature where the spatial coordinate of the pedestrian is embedded through a single layer with ReLU non-linearity. In the following equation, this layer is denoted as $\phi_e(.)$ with W_e being its embedding weights.

$$e_t^i = \phi_e(\Delta X_t^i; W_e) \qquad (1)$$

The spatial input will be provided as displacements $\Delta X_t^i = (x_t^i - x_{t-1}^i, y_t^i - y_{t-1}^i)$ which is the relative distance between pedestrian i's current time step t and the previous time step $t-1$.


Fig. 1. The overall framework of our proposed pedestrian trajectory prediction. The embedded spatial feature, the pedestrian-pedestrian interaction and the pedestrian-vehicle interaction features are concatenated before being inputted into the LSTM module.

2) Interaction embedded features: How interactions with other agents (i.e., pedestrians and vehicles) can influence a pedestrian's future trajectory is captured through the interaction feature input. These interactions are captured through a polar collision grid map in form of a tensor that will be discussed in more detail in Section III-D. We have separate tensors for encoding pedestrian-pedestrian (PPCG) and pedestrian-vehicle interaction (VPCG), as a vehicle can influence a pedestrian's trajectory differently from how another neighbouring pedestrian does. Each of these two tensors will be processed by going through an embedding layer with ReLU nonlinearity (ϕ_p for ped-ped and ϕ_m for ped-veh[1] interaction) before being used as an input to the LSTM network.

$$p_t^i = \phi_p(PPCG_t^i; W_p) \qquad (2)$$

$$m_t^i = \phi_m(VPCG_t^i; W_m) \qquad (3)$$

3) Prediction module: The three inputs of spatial feature (e_t^i), ped-ped interaction feature (p_t^i) and ped-veh interaction feature (m_t^i) will be concatenated and used as input to the LSTM cell with W_l weights.

$$h_t^i = LSTM(h_{t-1}^i, concat(e_t^i, p_t^i, m_t^i); W_l) \qquad (4)$$

The output of the LSTM will then go through a linear layer with W_o embedding weight to output the parameters of the bi-variant Gaussian distribution. The predicted position can then be generated by sampling from this distribution.

$$\hat{Y}_t^i = [\mu_x^i, \mu_y^i, \sigma_x^i, \sigma_y^i, \rho^i]_t = W_o h_t^i \qquad (5)$$

$$\tilde{Y}_{t+1}^i \sim \mathcal{N}(\hat{Y}_t^i) \qquad (6)$$

4) Loss function: Given that the model is predicting the parameters of a bi-variant Gaussian distribution ($\hat{Y}_t^i$) over the next positions, a negative log-likelihood loss function is used for learning the parameters of the model as follows:

$$Loss = -\sum_{i,t} log(P(Y_{t+1}^i | \tilde{Y}_{t+1}^i)) \qquad (7)$$

[1]Ped stands for pedestrians. Veh stands for vehicle

C. Interacting Agents

In this paper, we rely on collision risks for indicating each target pedestrian's nearby interacting agents. Only these specified interacting agents will then be considered for constructing the interaction feature of each target pedestrian. This method is opposed to other approaches followed in the literature that consider either all the agents present in the scene or rely only on distance for specifying interacting agents and building the interaction feature. The intuition is that at each time step a pedestrian's trajectory will be affected only by those agents that are in a collision course with it if the pedestrian keeps its current trajectory trend.

We use the feature of Time To Collision (TTC) as a proxy for measuring the collision risk with others. Therefore, at each time step, having the position and velocity vectors we calculate the time left for the agents to get closer to each other than a minimum comfortable distance (d_{min}), which we define as collision. This time between agent i and j with position vectors (x, y) and velocity vectors (v_x, v_y) is calculated according to eq 8. The notations used are for the information at time step t but we have removed the subscript t for simplicity.

$$TTC^{ij} = \tag{8}$$
$$\frac{-(\vec{D}_{rel} \cdot \vec{V}_{rel}) \pm \sqrt{(\vec{D}_{rel} \cdot \vec{V}_{rel})^2 - |\vec{V}_{rel}|^2(|\vec{D}_{rel}|^2 - d_{min}^2)}}{|\vec{V}_{rel}|^2}$$

In the above equation, the vectors for relative distance and the relative velocity between agents i and j are denoted as $\vec{D}_{rel} = X^i - X^j$ and $\vec{V}_{rel} = V^i - V^j$ respectively. Moreover, (.) denotes the dot product of two vectors and $|\vec{D}|$ is the size of the vector $\vec{D}$.

Using equation 8, all the agents with a valid TTC of greater than zero and less than a specified threshold of TTC_{thre} will be considered for constructing the interaction features as discussed in the next section. For any agent that might already be within d_{min} distance of the target pedestrian a TTC of zero will be considered.

D. Polar Collision Grid Interaction Feature

At each time step, having the specified interacting agent and their TTC, we construct the interaction feature in form of a polar grid centred at each target pedestrian's current position and in the direction of its current velocity vector.

We follow a similar approach to how we relied on collision risk for specifying interacting agents for building the interaction effect feature.

With the justification that a pedestrian follows relatively similar evasive manoeuvres for agents in the same collision course in terms of approaching angle, we consider the angle between the velocity vectors of two potentially colliding agents as another feature for encoding interaction and learning the trend of the evasive behaviour. Therefore, we discretize the angular space around each pedestrian's current velocity vector to a specified number of sectors (n_{sector}) and for each conflicting agent, we calculate the discretized sector where its approach angle lies on. The angles are calculated
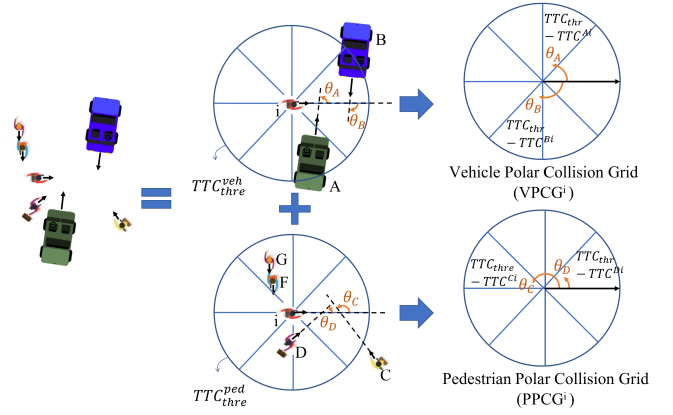

Fig. 2. The process of building the interaction features for pedestrian i according to the neighbouring agent's time to collision (TTC) and approach direction. Eight discrete sectors around the target pedestrian are considered in this figure for the approach angles. Only neighbouring agents with a TTC lower than a threshold will be considered in the construction of the interaction feature. The interaction feature for neighbouring vehicles and pedestrians are constructed separately.

with respect to the target pedestrian's current velocity vector in a counter-clockwise direction. This process results in a polar grid map around each pedestrian's current velocity vector while specifying the presence of each interacting agent in each cell of this polar grid.

To convert this polar grid map into an interaction feature that also includes collision risk information, we store the $TTC_{thr} - TTC$ value of each interacting agent in its corresponding cell. We aggregate the information of the multiple interacting agents by focusing on the riskiest agent in each approach angle cell which is the agent with the lowest time to collision or the highest value of $TTC_{thr} - TTC$. The resulting grid will be called the Polar time-to-Collision Grid (PPCG) which is a tensor of size $1 \times n_{sector}$ calculated according to equation 9 for each pedestrian i at time step t. In this equation, $\mathbb{1}_n[\theta]$ is a binary indicator function to specify if θ (the angle between vector $\vec{V}_t^i$ and $\vec{V}_t^j$) is in the n^{th} cell of the grid, and N^i is the set of all interacting agent of pedestrian i.

$$PPCG_t^i(n) = \max_{j \in N^i}(\mathbb{1}_n[\angle(\vec{V}_t^i, \vec{V}_t^j)](TTC_{thr} - TTC^{ij})) \tag{9}$$

This tensor will be constructed separately for pedestrians (PPCG) and vehicles (VPCG) that are interacting with a pedestrian. Fig. 2 visually shows the process of creating these interaction features for pedestrian i in the middle for a sample scenario. In this figure pedestrians F and G, despite being close to pedestrian i, are not included in the interaction feature since they are not causing any collision risks.

These interaction features will then be embedded according to equations 2 and 3 and used as an input to the LSTM.

IV. EXPERIMENT

We have evaluated the effectiveness of our proposed method by first comparing the overall methods with a couple

of baseline methods in the literature. Then we focus on the effectiveness of each of our introduced modules by conducting oblation studies.

A. Dataset

Since our method studies the effect of both pedestrian-pedestrian and pedestrian-vehicle interaction on the trajectory prediction of pedestrian agents, we require a dataset that contains both pedestrians and vehicles for our training and evaluation process. Meanwhile, we are willing to train our model in an environment where the interactions are less regulated by traffic rules and more governed through negotiation and social etiquette as these interaction patterns are more complicated to model compared to behaviours forced by traffic rules or emerged from strict environmental structures (e.g, crosswalks). Therefore, datasets from shared space environments fit our objectives well since traffic rules and road markings are alleviated in these environments and all different types of agents are encouraged to share the same space [27]. Here we use the HBS dataset [28] where its real-world trajectory data has been captured from a shared space close to a busy train station in Germany from a bird's eye view. This dataset has already been used in other trajectory prediction methods in the literature [9], [6], [29]. There are 338 vehicles, 1115 pedestrians, and 22 cyclists in the 30-minute video collected from this area and the trajectory data are available at 2 Hz. The dominant number of pedestrians in this dataset also suits our method which is focused on predicting only pedestrians' trajectories. Same as [9], [28], [30], we use the first 10 minutes (31% of the dataset) as the test set and the remainder for training. Here we observe 6 time-steps covering 3 seconds of the trajectory and predict the next 6 time steps (3 seconds) of the trajectory, similar to [9], [30].

B. Evaluation Metrics and Baselines

In line with prior works we use the following metrics to evaluate the method and report the prediction errors.

- Average displacement error (ADE): The Euclidean distance between predicted and ground truth positions at each time step averaged over all predicted time steps.
- Final Displacement error (FDE): The Euclidean distance between predicted and ground truth positions at the final predicted time step.
- Modified Hausdorff Distance (MHD): The largest distance between any two points on predicted and ground truth trajectories (without considering time step alignment).
- Speed error (SE): The root mean square error (RMSE) between the speed of the predicted trajectory and the ground truth speed at each time step.
- Heading error (HE): The RMSE between the heading of the predicted trajectory and the ground truth heading at each time step.

We compare the performance of our method against the following baseline methods:

- Linear Regression (LR): A linear regression model for each position dimension individually.
- Vanilla LSTM: A naive LSTM considering only trajectory history without the interaction effects from pedestrians and vehicles.
- Social LSTM: The LSTM model with the social pooling layer proposed by Alahi et al. [1] considering only pedestrian-pedestrian interaction.

The grid used in the Social LSTM model as the interaction encoding module is build based on relative position data. Thus, it is used as a baseline for evaluating our proposed interaction grid constructed using TTC information.

C. Implementation Details

The embedding layer for the spatial inputs and the two interaction feature all have a size of 64. We used a hidden state dimension of 128 for the LSTM module. The number of sectors (n_{sector}) for the polar collision grid was set to 8. A TTC_{thre} of 9 seconds with a d_{min} of 0.7 meters was used for the pedestrian-pedestrian interaction features. These values were 8 seconds and 1 meter for the pedestrian-vehicle interaction feature respectively. These values were selected after further tuning within a reasonable range according to the physical dimensions of the agents and the common personal space of humans [31] as well as the agents' minimum distance and speed ranges in the HBS dataset. Other hyperparameters of the model were also tuned by testing different values and selecting the one that resulted in the lowest error. A training batch size of 10 was used for 200 epochs. We used the RMS-prop optimizer with a learning rate of 0.001. The Nvidia GeForce RTX 2080 Ti GPU was used for the training and testing processes.

V. RESULTS AND DISCUSSION

A. Quantitative Evaluation

The performance of the trained models on the HBS test set are summarized in Table I Since the output of the deep-learning models is a distribution over the predicted position, for evaluation, same as [7] we generated 20 samples for each model and reported the best result. Lower numbers are better. The best values among the different models for each metric are shown in boldface. The results show that our proposed method outperforms the baseline methods on all the evaluation metrics except for heading angle error. These results demonstrate the advantage of formulating interaction according to collision probability and approaching directions instead of only relying on relative distances. According to Table I, Social LSTM has slightly higher prediction errors than the Vanilla LSTM which does not consider interactions. This is consistent with the result reported by others for these two models using other datasets [2], [7].

Oblation study: To more clearly show the benefit of our proposed interaction modules for both pedestrian-pedestrian and pedestrian-vehicle interaction, we conducted an oblation study where we took out the interaction modules in a step-wise process and compared prediction errors with and without those modules in place.

Models \ Metrics	ADE (m)	FDE (m)	MHD (m)	SE (m/s)	HE (°)
Linear regression	0.696	1.238	2.995	0.390	44.0
Vanilla LSTM	0.305	0.676	2.855	0.240	32.7
Social LSTM	0.309	0.677	2.852	0.244	**31.5**
P-CollisionGrid	0.304	0.664	2.811	0.235	32.6
V-CollisionGrid	0.305	0.669	2.827	**0.232**	32.9
PV-CollisionGrid	**0.295**	**0.648**	**2.791**	0.235	31.7

Models \ Metrics	ADE (m)	FDE (m)	MHD (m)	SE (m/s)	HE (°)
Social LSTM	0.309	0.677	2.852	0.244	**31.5**
Social LSTM +Filtered interaction	**0.298**	**0.658**	**2.827**	**0.234**	32.3

Following this, we have trained our model by taking out the pedestrian-vehicle interaction feature and only keeping the collision grid associated with the pedestrian-pedestrian interaction. We called this the *P-CollisionGrid*. We repeated the same process by this time taking out the pedestrian-pedestrian interaction feature and only keeping the pedestrian-vehicle collision grid and called the model *V-CollisionGrid*. The prediction errors of these two models along with our complete model *PV-CollisionGrid* are also reported in Tabel I.

The results show that adding each of the interaction modules positively improves the prediction accuracy of the model in terms of ADE, FDE, MHD and heading error. This indicates the effectiveness of our interaction module in capturing the influence of potentially conflicting agents on the predicted trajectory of a pedestrian. However, the lowest speed error is for the model that only considers pedestrian-vehicle interaction. This could be due to the high influence that vehicles have on pedestrians' speed adjustment along their trajectory especially when a pedestrian stops for giving way to a vehicle.

Interacting agent selection: For more specifically analyzing the effect of our proposed method for the selection of interacting agents based on TTC, we compared two versions of the Social LSTM baseline model. The first version is the original one where the interaction feature is constructed using agents closer than a specified distance. In the second version, the neighbouring agents are first filtered based on our proposed TTC criteria and then the original social pooling module [1] is constructed considering those agents with collision risks.

The performance evaluation of these two versions in Table II shows that filtering the agent to consider only the ones in a conflict course with the target agent will actually improve the prediction accuracy. According to this result, removing the information of those agents that do not produce a collision risk on the target agent seems to be not only harmless but also beneficial as the target agent's trajectory will not be affected by those agents. Therefore, our proposed method of selecting interacting agents based on collision risk seems to outperform the distance-based methods by concentrating on those agents that actually make the target agent adjust its path for avoiding possible collisions.

B. Qualitative Evaluation

To qualitatively analyze the predicted trajectory of our model compared to other baselines, we plotted the prediction outputs for a couple of interactive scenarios in the test set as shown in Fig. 3-5.

In the scenario shown in Fig. 3, our method detects vehicle B as an interacting agent with pedestrian A in the last observed timestep (t_{obs}), meaning that a collision will take place if the two agents continue their velocity as it is at time t_{obs}. Having identified this interaction, our method then considers Vehicle B's information on TTC and approach angle when predicting pedestrian A's future trajectory. Therefore, as shown in Fig. 3, our method predicts an avoidance manoeuvre which is close to the ground truth. This is while the Social SLTM method which does not consider the pedestrian-vehicle interaction, predicts a trajectory that ends up in a collision with the vehicle at the last time step of the prediction window. This supports the effectiveness of our pedestrian-vehicle interaction module.

Fig. 5 shows another example focused on pedestrian-pedestrian interactions and compares the prediction of our Collision Grid method with the predictions of the Social LSTM method that also models the interaction between pedestrians. Focusing on Pedestrian C as a target agent which is indicated with a star marker at time step t_{obs}, we show all its neighbours at the same time step with a square marker. In each of the methods, the neighbouring agents that will be considered for constructing the interaction feature are
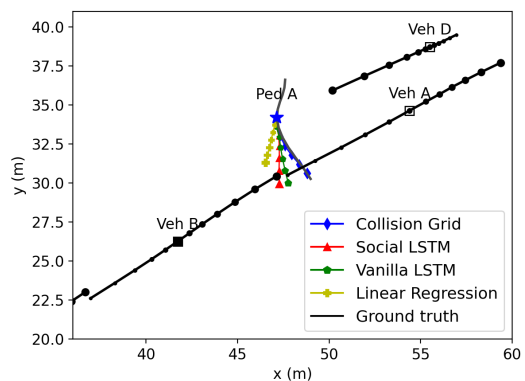

Fig. 3. The predicted trajectories for our method compared to the baseline methods for a pedestrian-vehicle interaction scenario in the HBS test dataset. A marker with a larger size represents a later time step along the trajectory. The position of the vehicles at the last time step of the observation period is indicated with a square marker. Our method detects the target pedestrian A (indicated with start) to have interaction with vehicle B at this time step according to the TTC criteria. This identified interacting vehicle is indicated with a filled square marker at this time step. Considering this interacting vehicle our method has predicted a trajectory that avoids the predicted collision while being close to the ground truth trajectory.
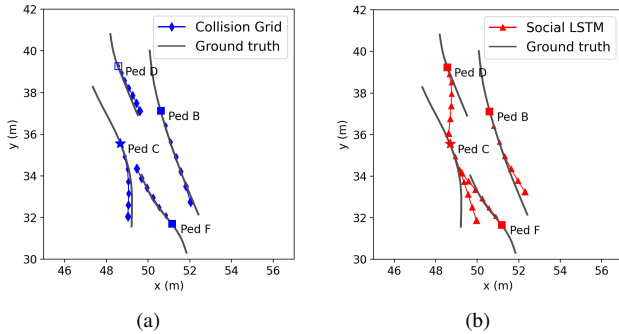
Fig. 5. Prediction of our Collision Grid method compared to Social LSTM for a pedestrian-pedestrian interaction scenario. Our method produces a prediction closer to the ground truth for the pedestrians such as Ped C. This is done by detecting Ped C's interacting agents (Ped F and Ped B) more accurately while relying on time to collision for interaction modelling as opposed to the Social LSTM model that considers all agents closer to the target pedestrian than a threshold distance regardless of their moving direction which happens to be all the agents present in the scene here. Interacting agents for Ped C (target) are indicated with a filled square marker at time step t_{obs} while markers for non-interacting agents are unfilled.

indicated with a filled square marker. According to Fig. 5 while the Social LSTM considers all the agents present in this scene in building the interaction, our method only detects pedestrian F and B as interacting agents for pedestrian C and therefore relies only on these pedestrians' information for constructing the interaction feature. This results in a more accurate prediction for Pedestrian C's future trajectory with less deviation from the ground truth compared to the Social LSTM method. This shows the effectiveness of our proposed interacting agent selection process that focuses specifically on neighbouring agents that can actually affect the target agent's trajectory due to being in a potential collision course, as opposed to other methods such as Social LSTM that consider all agents that are closer than a distance to the target agent regardless of their moving direction. This latter approach can introduce extra not useful information for interaction consideration that can lead to an undesired prediction as seen in Fig. 5 for Social LSTM.

More qualitative examples of our method's prediction compared to other baseline methods can be found in Fig. 4 for other scenarios. There are cases where our method has more deviation from the ground truth compared to other methods. Fig. 4(c) provides an example in which two pedestrians (Ped A and Ped B) are walking together as a group. For this case our method has a false positive collision detection and predicts a trajectory for Ped B that is slightly to the right of the ground truth, getting away from Ped A, instead of predicting close parallel paths for these two pedestrians. However, despite the poor prediction in these particular cases, our algorithm produces less error on average compared to other baselines' predictions as seen in Table I. In future, the prediction of our algorithm can be further improved by adding a group detection for filtering out possible false collision detection in the model.

VI. LIMITATION AND FUTURE WORK

Our work had several limitations. The goal was to model the interaction effect between pedestrians and vehicles to predict future trajectories of pedestrians. These complicated interactions mostly happen in environments that do not segregate the operational space of the pedestrians and vehicles, e.g., shared spaces. Therefore, we limited our datasets to shared space environments that greatly embed these high-interactive behaviours. However, only a few datasets exist from these shared spaces while in some of them the number of data points and trajectories are limited. Thus, we only trained and tested our algorithm on one dataset (HBs) which outperformed the baselines. Future work can benefit from testing on other datasets to confirm whether the results can be generalized to other environments. Also, for focusing on patterns that could be derived from only trajectory data, we did not consider the effect of the environment's structure. Future work can take the map of the environment as another input to study the effects of these other factors.

VII. CONCLUSION

We proposed the use of collision-risk information for encoding the interaction effects for pedestrian trajectory prediction in an environment shared with vehicles. Interacting pedestrians and vehicles with a target pedestrian were first
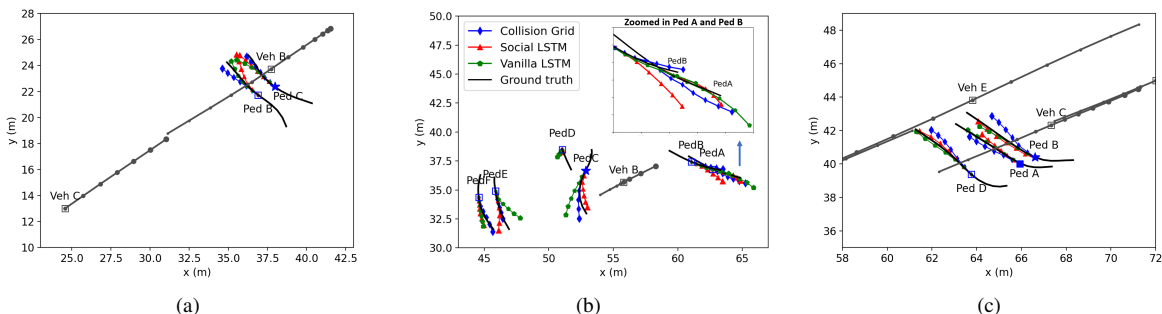

Fig. 4. The predicted trajectories for our method compared to the deep-learning-based baseline methods in the HBS dataset. Our predictions overall are closer to the ground truth compared to the baselines. Case (c) is an example with a non-accurate prediction from our *Collision Grid* method where a probable group interaction between Ped A and Ped B is mistakenly detected as a possible collision ending up with a prediction for Ped B that deviates to the right of the ground truth data. Interacting agents with a sample target pedestrian (shown with star) are shown with a filled square marker.

selected based on the heuristic of Time-to-Collision (TTC). A novel polar collision grid map for each target pedestrian then encoded the interaction effects using the information of TTC and the approach direction of these interacting agents.

Our results showed that by filtering the neighbouring agents for focusing only on those agents that could actually affect the target pedestrians through posing a collision risk, we may better capture the interaction effects on the predicted trajectories. Also, by directly using the information on the approaching risk direction and the criticality of the situation using TTC, we were able to predict the corresponding changes in the trajectory more accurately than the baseline methods. According to the results, we believe that providing the prediction network with more relevant engineered features can guide deep-learning-based prediction methods to better learn the patterns underlying the trajectories, leading to predictions that are closer to the ground truth.

REFERENCES

[1] 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*, 2016, pp. 961–971.

[2] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2255–2264.

[3] H. Xue, D. Q. Huynh, and M. Reynolds, "Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction," in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 1186–1194.

[4] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*. Springer, 2020, pp. 683–700.

[5] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 336–345.

[6] H. Cheng, W. Liao, M. Y. Yang, M. Sester, and B. Rosenhahn, "Mcenet: multi-context encoder network for homogeneous agent trajectory prediction in mixed traffic," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–8.

[7] C. Zhang and C. Berger, "Learning the pedestrian-vehicle interaction for pedestrian trajectory prediction," in *2022 8th International Conference on Control, Automation and Robotics (ICCAR)*. IEEE, 2022, pp. 230–236.

[8] S. Eiffert, K. Li, M. Shan, S. Worrall, S. Sukkarieh, and E. Nebot, "Probabilistic crowd gan: Multimodal pedestrian trajectory prediction using a graph vehicle-pedestrian attention network," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5026–5033, 2020.

[9] H. Cheng and M. Sester, "Modeling mixed traffic in shared space using lstm with probability density mapping," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 3898–3904.

[10] Y. Hu, S. Chen, Y. Zhang, and X. Gu, "Collaborative motion prediction via neural motion message passing," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 6319–6328.

[11] X. Mo, Z. Huang, Y. Xing, and C. Lv, "Multi-agent trajectory prediction with heterogeneous edge-enhanced graph attention network," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 9554–9567, 2022.

[12] J. Li, H. Ma, Z. Zhang, J. Li, and M. Tomizuka, "Spatio-temporal graph dual-attention network for multi-agent prediction and tracking," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10 556–10 569, 2021.

[13] M. Golchoubian, M. Ghafurian, K. Dautenhahn, and N. L. Azad, "Pedestrian trajectory prediction in pedestrian-vehicle mixed environments: A systematic review," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–24, 2023.

[14] H. Bi, Z. Fang, T. Mao, Z. Wang, and Z. Deng, "Joint prediction for kinematic trajectories in vehicle-pedestrian-mixed scenes," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 10 383–10 392.

[15] S. Carrasco, D. F. Llorca, and M. Sotelo, "Scout: Socially-consistent and understandable graph attention network for trajectory prediction of vehicles and vrus," in *2021 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2021, pp. 1501–1508.

[16] H. Girase, H. Gang, S. Malla, J. Li, A. Kanehara, K. Mangalam, and C. Choi, "Loki: Long term and key intentions for trajectory prediction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9803–9812.

[17] W.-C. Lai, Z.-X. Xia, H.-S. Lin, L.-F. Hsu, H.-H. Shuai, I.-H. Jhuo, and W.-H. Cheng, "Trajectory prediction in heterogeneous environment via attended ecology embedding," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 202–210.

[18] M. Prédhumeau, L. Mancheva, J. Dugdale, and A. Spalanzani, "Agent-based modeling for predicting pedestrian trajectories around an autonomous vehicle," *Journal of Artificial Intelligence Research*, vol. 73, pp. 1385–1433, 2022.

[19] M. Predhumeau, L. Mancheva, J. Dugdale, and A. Spalanzani, "An agent-based model to predict pedestrians trajectories with an autonomous vehicle in shared spaces," in *AAMAS 2021-20th International Conference on Autonomous Agents and Multiagent Systems*, 2021, pp. 1–9.

[20] J. Cheng and D. Li, "A study of interference between pedestrians and vehicles in drop-off area at railway station based on celluar automata," in *2019 5th International Conference on Transportation Information and Safety (ICTIS)*. IEEE, 2019, pp. 913–919.

[21] X. Li and J.-Q. Sun, "Studies of vehicle lane-changing to avoid pedestrians with cellular automata," *Physica A: Statistical Mechanics and its Applications*, vol. 438, pp. 251–271, 2015.

[22] Z. Zuo, X. Zhang, and Y. Wang, "Map feature based trajectory prediction with multi-class traffic participants," in *2021 40th Chinese Control Conference (CCC)*. IEEE, 2021, pp. 7312–7317.

[23] M. Huber, Y.-H. Su, M. Krüger, K. Faschian, S. Glasauer, and J. Hermsdörfer, "Adjustments of speed and path when avoiding collisions with another pedestrian," *PloS one*, vol. 9, no. 2, p. e89589, 2014.

[24] A. Vemula, K. Muelling, and J. Oh, "Social attention: Modeling attention in human crowds," in *2018 IEEE international Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4601–4607.

[25] J. Li, F. Yang, H. Ma, S. Malla, M. Tomizuka, and C. Choi, "Rain: Reinforced hybrid attention inference network for motion forecasting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16 096–16 106.

[26] Y. Su, J. Du, Y. Li, X. Li, R. Liang, Z. Hua, and J. Zhou, "Trajectory forecasting based on prior-aware directed graph convolutional neural network," *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[27] M. Prédhumeau, A. Spalanzani, and J. Dugdale, "Pedestrian behavior in shared spaces with autonomous vehicles: An integrated framework and review," *IEEE Transactions on Intelligent Vehicles*, 2021.

[28] F. Pascucci, N. Rinke, C. Schiermeyer, V. Berkhahn, and B. Friedrich, "A discrete choice model for solving conflict situations between pedestrians and vehicles in shared space," *arXiv preprint arXiv:1709.09412*, 2017.

[29] F. T. Johora and J. P. Muller, "On transferability and calibration of pedestrian and car motion models in shared spaces," *Transportation letters*, vol. 13, no. 3, pp. 172–182, 2021.

[30] H. Cheng and M. Sester, "Mixed traffic trajectory prediction using lstm–based models in shared space," in *Geospatial Technologies for All: Selected Papers of the 21st AGILE Conference on Geographic Information Science 21*. Springer, 2018, pp. 309–325.

[31] A. Gorrini, K. Shimura, S. Bandini, K. Ohtsuka, and K. Nishinari, "Experimental investigation of pedestrian personal space: toward modeling and simulation of pedestrian crowd dynamics," *Transportation research record*, vol. 2421, no. 1, pp. 57–63, 2014.