

Probabilistic Crowd GAN: Multimodal Pedestrian Trajectory Prediction using a Graph Vehicle-Pedestrian Attention Network

Stuart Eiffert*, Kunming Li*, Mao Shan, Stewart Worrall, Salah Sukkarieh and Eduardo Nebot

Abstract—Understanding and predicting the intention of pedestrians is essential to enable autonomous vehicles and mobile robots to navigate crowds. This problem becomes increasingly complex when we consider the uncertainty and multimodality of pedestrian motion, as well as the implicit interactions between members of a crowd, including any response to a vehicle. Our approach, Probabilistic Crowd GAN, extends recent work in trajectory prediction, combining Recurrent Neural Networks (RNNs) with Mixture Density Networks (MDNs) to output probabilistic multimodal predictions, from which likely modal paths are found and used for adversarial training. We also propose the use of Graph Vehicle-Pedestrian Attention Network (GVAT), which models social interactions and allows input of a shared vehicle feature, showing that inclusion of this module leads to improved trajectory prediction both with and without the presence of a vehicle. Through evaluation on various datasets, we demonstrate improvements on the existing state of the art methods for trajectory prediction and illustrate how the true multimodal and uncertain nature of crowd interactions can be directly modelled.

I. INTRODUCTION

Pedestrian motion prediction is required for the safe and efficient operation of autonomous vehicles and mobile robots in shared pedestrian environments, such as malls and campuses, as shown in Fig. 1. The motion of individual members of a crowd is dependent on the motion of others nearby, including any vehicles, and contains significant uncertainty during interactions. In order to better predict pedestrian motion we need to be able to model this uncertainty, which is often multimodal due to the variety of ways in which individuals can interact and avoid each other. Recent works that aim to capture this multimodal and probabilistic nature of crowd interactions have attempted to do so through repeated sampling of generative models, often using Recurrent Neural Networks (RNNs)-based autoencoders trained as Generative Adversarial Networks (GANs) [2], [3]. Due to the nature of adversarial training, where generated trajectories must match the form of the ground truth for comparison by the Discriminator, these methods are limited to generating non-probabilistic outputs. Instead, they require repeated sampling with use of a random latent variable to identify the true multimodal distribution during inference.

Additionally, in applications involving the use of a single-vehicle around pedestrians, such as autonomously navigating a university campus, accurate prediction of nearby pedestrian motion requires inclusion of vehicle-pedestrian interactions in any predictive model. Recent work [3] has shown that the use of Graph Attention Networks (GATs) [4] can improve

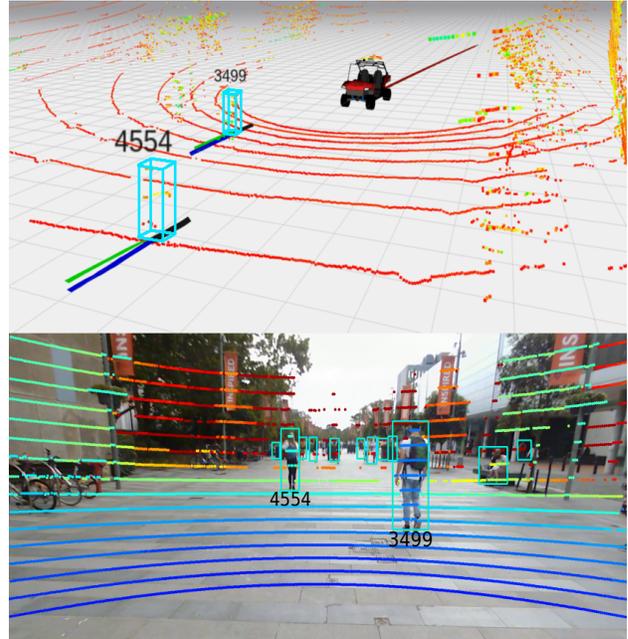


Fig. 1: Motion of detected pedestrians is predicted using our method Probabilistic Crowd GAN with a Graph Vehicle-Pedestrian Attention Network (PCGAN). Observed trajectories are shown in black. The most likely modal path of the multimodal probabilistic prediction are shown in green against ground truth in blue. Predictions use the vehicle’s motion as a feature input on the USyd Campus dataset [1].

the modelling of social interactions between pedestrians, as compared to previously used social pooling layers.

Our proposed method, Probabilistic Crowd GAN (PCGAN), allows for the direct prediction of probabilistic multimodal outputs during adversarial training. We make use of a Mixture Density Network (MDN) within the GAN’s generator to output a Gaussian mixture model (GMM) for each pedestrian, demonstrating how clustering of each component of the GMM allows the finding of likely modal paths, that can then be compared to ground truth trajectories by the GAN’s discriminator. Additionally, we extend the use of GATs for modelling crowd interactions to include heterogeneous interactions between a vehicle and pedestrians, as a Graph Vehicle-Pedestrian Attention Network (GVAT) used for modelling social interactions in our method. We validate our approach on several publicly available real world datasets of pedestrian crowds, as well as two datasets which include crowd-vehicle interactions.

*The first two authors contributed equally to this work

Our main contributions in this work include:

- Direct multimodal probabilistic output from a GAN for trajectory prediction.
- Extension of Graph Attention Networks to include a shared vehicle feature in the pooling mechanism.
- Improved pedestrian motion prediction both with and without the presence of a single vehicle.

II. RELATED WORK

Pedestrian Trajectory Prediction: Approaches to motion prediction in crowds have tended to focus either on modelling scene-specific motion patterns through the inclusion of contextual features [5] and learning crowd motion for a specific observed scene [6], [7], or on interactions between individuals. Crowd interactions have been modelled using either hand-crafted methods such as the Social Force Model (SFM) [8], or using learnt models of interaction. Recent developments in learning-based methods of trajectory prediction such as RNNs [9], [10] allow for improved prediction in crowded environments, outperforming parametric based methods such as SFM [11]. These methods have been applied to multimodal prediction by learning semantically meaningful latent representations in conditional variational autoencoders [12], [13] and GANs [14], or through clustering modal paths in output distributions [15]. However, these methods can still fail to outperform even simple baselines such as constant velocity models in many situations [16].

GANs for Probabilistic Prediction: GANs [17] have been recently used to enable the generation of socially acceptable trajectories in crowd motion prediction. Gupta et al. [2] proposed Social GAN, in which the generator of the network consists of an LSTM based encoder-decoder with a social pooling layer modelling the relationship between each pedestrian. The output trajectories of the generator are directly compared to the ground truth by the Discriminator. Social-BiGAT [3] extends this idea, introducing a flexible graph attention network and further encouraging generalization towards a multi-modal distribution. This method, as well as a similar GAN based approach proposed by [18], also make use of overhead contextual scene inputs, which are often difficult to capture in autonomous driving systems. Prior work [2], [5], [14] using GANs for trajectory prediction has followed the assumption from GAN application to image synthesis that we cannot efficiently evaluate the output distribution, but can sample from it, requiring multiple iterations to identify the true multimodal distribution. However, our problem’s output distribution is much lower-dimensional than image synthesis, and has been modelled previously by GMMs as in [12], [15], [19], allowing a distribution to be generated from a single iteration. Further, our aim differs from the synthesis in that we are not trying to just generate samples in the style of ground truth conditioned on an observation, but rather samples that mimic ground truth.

Interaction Modelling: Alahi et al. [9] proposed the use of RNNs with a social pooling layer to capture interactions between pedestrians in a crowd, with similar pooling layer being used in Social GAN [2]. The pooling mechanism used

in [20] allows interactions between different agent types by learning respective weightings for each relationship. Recent works [21], [22] have extended the work of Vemula et al. [10] to apply Structural-RNNs [23] to heterogeneous interactions, modelling multiple road agents types using RNNs in a spatio-temporal graph. Veličković proposed graph attention networks (GAT) [4] to implicitly assign different importance to nodes in graph structured data. Kosaraju et al. [3] applied this concept to multimodal trajectory prediction by formulating pedestrian interactions as a graph, however apply the graph structure only within the pooling mechanism as a GAT, rather than modelling each relationship of the graph as a separate RNN as in [10]. We extend this idea in our work to demonstrate how a vehicle’s feature can be included in GAT as a GVAT.

III. METHOD

A. Problem Definition

In this paper, we address the problem of pedestrian trajectory prediction in crowds both with and without the presence of a vehicle. Given observed trajectories \mathbf{X} , and the vehicle path \mathbf{V} , for all time steps in period $t \leq T_{obs}$, where $\mathbf{X}^t = [\mathbf{X}_1^t, \mathbf{X}_2^t, \dots, \mathbf{X}_N^t]$ for N pedestrians within a scene, our aim is to predict the likely future trajectories $\bar{\mathbf{Y}}^t = [\bar{\mathbf{Y}}_1^t, \bar{\mathbf{Y}}_2^t, \dots, \bar{\mathbf{Y}}_N^t]$ for each pedestrian in N , across a future time period $T_{obs} < t \leq T_{pred}$. The input position of the i th pedestrian at time t is defined as $\mathbf{X}_i^t = (x_i^t, y_i^t)$ and the vehicle as $\mathbf{V}^t = (x_v^t, y_v^t)$. We denote \mathbf{Y} as the ground truth future trajectory, with the position of the i th pedestrian at time t defined as $\mathbf{Y}_i^t = (x_i^t, y_i^t)$ and predicted position as $\bar{\mathbf{Y}}_i^t = \{(x_i^t, y_i^t, \bar{w}_i^m)_{m=1}^M\}$ for all predicted modal paths $m \in M$, where \bar{w}_i^m is the likelihood of the predicted modal path m for agent i . $\bar{\mathbf{Y}}_i^t$ is found from the probabilistic output $\hat{\mathbf{Y}}_i^t$, a Gaussian mixture model (GMM) detailed in Eq. 14.

B. Overview

Our approach consists of two networks, a Generator and a Discriminator trained adversarially. The Generator is composed of an RNN encoder, our GVAT module, an RNN decoder, and a Mixture Density Network (MDN). The Discriminator is composed of the modal path clustering module MultiPAC, an RNN encoder and a multilayer perceptron (MLP). Fig. 2 illustrates the overall system architecture. **Generator:** The Generator is based on an RNN encoder-decoder framework using LSTM modules, where the GVAT Pooling module is applied to the hidden states between the encoder and decoder LSTMs. The input to the encoder LSTM at each timestep $t < T_{obs}$ is the observed position of each pedestrian $i \in N$, which is first passed through a linear embedding layer ϕ_e as follows:

$$e_i^t = \phi_e(x_i^t, y_i^t; W_{emb}^e) \quad (1)$$

$$h_{ei}^t = LSTM(h_{ei}^{t-1}, e_i^t; W_{enc}) \quad (2)$$

where W_{emb} is embedding weight of ϕ_e . All pedestrians within a scene share the LSTM weights W_{enc} . The decoder’s initial hidden state at $t = T_{obs}$ is composed of the encoder’s

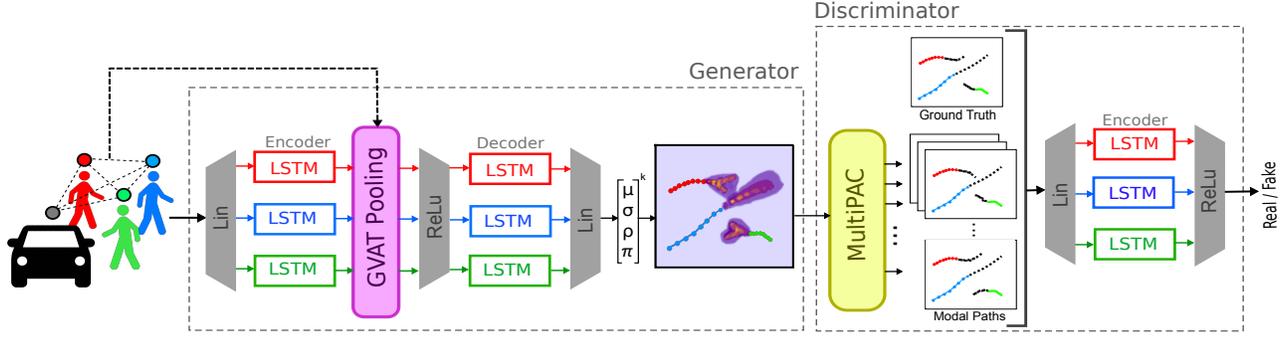


Fig. 2: Observed pedestrian trajectories are passed to the Generator’s encoder LSTM, whilst the relative position of all agents, including any vehicle, are passed to the GVAT Pooling module. The Generator outputs a GMM for each agent, from which the MultiPAC module finds the likely modal paths, which are compared to ground truth paths by the Discriminator.

final hidden state, concatenated with the transformed output of GVAT Pooling for each agent, detailed further in *Section III-C*. The first input to the decoder at $t = T_{obs}$ is again the observed pedestrian positions, passed first through a linear embedding layer ϕ_d in the same form as ϕ_e with separate weights. However, as the decoder outputs are a distribution, rather than a single point, we do not simply pass the prediction from the prior timestep as input to the decoder’s current timestep. Instead, for all prediction timesteps $T_{obs} < t \leq T_{pred}$ the decoder inputs are zeros. This is done as opposed to other probabilistic approaches which feed a sample from the prior output as current input to the decoder. This *zero-feed* approach is performed for both training and inference, and has been shown to improve performance for probabilistic outputs [15].

$$g_i^t = GVAT(\mathbf{X}^t, \mathbf{V}^t, h_e^t) \quad (3)$$

$$h_{di}^t = LSTM(MLP_{dec}(h_{di}^{t-1}, g_i^t), d_i^t; W_{dec}) \quad (4)$$

$$d_i^t = \begin{cases} \phi_d(x_i^t, y_i^t; W_{emb}^d), & t \leq T_{obs} \\ 0, & t > T_{obs} \end{cases}$$

where $h_{di}^{T_{obs}} = h_{ei}^{T_{obs}}$ and h_e^t is the combined output of Eq.2 for all agents in the scene. MLP_{dec} is a multilayer perceptron with ReLU non-linearity and W_{dec} is the embedding weight. The outputs of the decoder are passed through a linear embedding layer ϕ_{mdn} with weights W_{mdn} that maps to a bivariate GMM output \hat{Y}_i^t for each agent’s position at each predicted timestep. \hat{Y}_i^t is then passed to the MultiPAC module to determine the set of likely modal paths \bar{Y}_i^t :

$$\hat{Y}_i^t = \phi_{mdn}(h_{di}^t; W_{mdn}) \quad (5)$$

Discriminator: The Discriminator is comprised of a MultiPAC module, and an LSTM encoder of the same form as the Generator’s, with separate weights. The output of Generator \hat{Y}_i^t is first passed to MultiPAC, from which we compute the set of likely modal paths \bar{Y}_i^t , as detailed in *Section III-E*. This produces trajectories in the same form as the ground truth, allowing comparison by the Discriminator’s

encoder. The encoder is applied across all timesteps $0 < t \leq T_{pred}$, with inputs first passed through a linear embedding layer. Outputs of the encoder are passed to a multilayer perceptron (MLP) with ReLU activation, classifying the path as either a Real or Fake.

Loss: Training of the network is achieved using two loss functions L_{lh} and L_{adv} . L_{lh} is the negative log-likelihood of the ground truth path \mathbf{Y} given the Generator G ’s output $\hat{\mathbf{Y}}$, across all prediction timesteps, for all pedestrians:

$$L_{lh} = - \sum_{t=T_{obs}+1}^{T_{pred}} \sum_i^N \log(P(Y_i^t | \hat{Y}_i^t)) \quad (6)$$

L_{adv} is the adversarial loss, determined from the binary cross entropy of the Discriminator D ’s classification of the modal paths $\bar{\mathbf{Y}}$ produced from $\hat{\mathbf{Y}}$ by MultiPAC:

$$L_{adv} = \mathbb{E}[\log(D(X_i, Y_i))] + \sum_{m=1}^M \mathbb{E}[\log(1 - w_i^m D(X_i, \bar{Y}_i))] \quad (7)$$

where the first term refers to D ’s estimate of the probability that the ground truth trajectory Y_i is real, and the second term is the sum of weighted estimates for each modal path in the set \bar{Y}_i being real. We combine the losses to find the optimal Discriminator D^* and Generator G^* , with weighting α applied to L_{lh} :

$$G^*, D^* = \underset{G}{\operatorname{argmin}} \underset{D}{\operatorname{argmax}} [L_{adv} + \alpha L_{lh}] \quad (8)$$

C. Graph Vehicle-Pedestrian Attention Network

We introduce a novel Graph Vehicle-Pedestrian Attention Network (GVAT), which extends upon the use of GATs [4] for trajectory prediction by [3] [24], allowing the modelling of social interactions between all pedestrians in a scene, and accommodating the inclusion of a vehicle if present. As opposed to [3], where only agent hidden states form the GAT input features, we also utilise distance between agents, so that vehicle distance to agent i can be included to allow the attention module to account for the impact that the vehicle’s motion has on each ped-ped relationship. *Fig. 3*. details the input features of a single node in the graph.

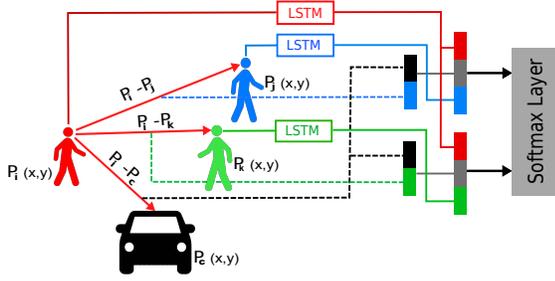


Fig. 3: Node features of agent i (red) in GVAT. The distance from i to the vehicle is appended to each other ped-ped distance input before encoding to account for the impact of the vehicle on i 's relationships within the graph. The input to softmax layer is $u_{i,j}^t$ as per Eq. 10

For the i th pedestrian, the input to the softmax layer is formulated across all other pedestrians $j \in N \setminus \{i\}$ by embedding the distance from pedestrian i to the neighbour pedestrian j and the vehicle. The softmax scalar $a_{i,j}^t$ is then used to scale the amount agent j 's hidden state influences agent i . The summed output across all other agents, g_i^t , is then concatenated with i 's original state to form the output of GVAT pooling h_{gi}^t . ϕ_r , ϕ_u and ϕ_{gat} are linear embedding functions, W_r , W_u and W_{gat} denote their parameters respectively:

$$r_{i,j}^t = \phi_r(x_i^t - x_j^t, y_i^t - y_j^t, z_i^t; W_r) \quad (9)$$

$$z_i^t = \begin{cases} (x_i^t - x_v^t, y_i^t - y_v^t), & \text{vehicle present} \\ (0, 0), & \text{no vehicle} \end{cases}$$

$$u_{i,j}^t = \phi_u(\text{concat}(r_{i,j}^t; h_{ej}^t; h_{ei}^t); W_u) \quad (10)$$

$$a_{i,j}^t = \frac{\exp(u_{i,j}^t)}{\sum_{k \in N \setminus \{i\}} \exp(u_{i,k}^t)} \quad (11)$$

$$g_i^t = \sum_{j \in N \setminus \{i\}} \phi_{gat}(a_{i,j}^t \cdot h_{ej}^t; W_{gat}) \quad (12)$$

$$h_{gi}^t = \text{concat}(h_i^t, g_i^t) \quad (13)$$

D. Mixture Density Network

An MDN is used to allow the Generator to propose a multimodal solution for each agent's future trajectory, with assigned relative likelihoods for each Gaussian component of the mixture model. To achieve this, the output of the Generator's decoder is passed through a multilayer perceptron (MLP) to produce output \hat{Y}^t in the form:

$$\hat{Y}^t = [\pi, \mu_x^k, \mu_y^k, \sigma_x^k, \sigma_y^k, \rho_{k=1}^{k=K}]^t \quad (14)$$

where K is the total number of components used in the mixture model, π is the weight of each component in the mixture, μ is the mean and σ the standard deviation per dimension, and ρ is the correlation coefficient, for each timestep $T_{obs} < t \leq T_{pred}$. This is performed separately for each agent $i \in N$, which has been left off for clarity.

E. Modal Path Clustering

In order to allow the training of the Discriminator, the output of the Generator must be converted to the same form as the ground truth trajectories \mathbf{Y} . This requires extracting individual tracks from the GMM $\hat{\mathbf{Y}}$, whilst preserving the multimodality of the distribution. We achieve this by adapting the multiple prediction adaptive clustering algorithm (MultiPAC) proposed by Zyner et al. [15] to allow backpropagation for use during training. MultiPAC finds the set of likely 'modal paths' $\bar{\mathbf{Y}}$, for each pedestrian from $\hat{\mathbf{Y}}$. It achieves this by clustering the components of the GMM at each timestep using DBSCAN [25], determining each cluster's centroid from the weighted average of all Gaussians in the mixture. Clusters in subsequent timesteps are assigned to parent clusters, forming a tree of possible paths with an upper limit of children at each timestep being the number of mixtures used within the GMM. This tree is computed from a single forward pass of the model, resulting in a forked trajectory when diverging possible paths are predicted for a single agent, passing each branch of the fork separately to the Discriminator. The paths from each leaf to the root are returned as the set of modal paths $\bar{\mathbf{Y}}$ for each pedestrian with assigned likelihoods w .

F. Implementation

The LSTM encoder and decoder of the Generator both have a hidden state size of 32, whilst the Discriminator's LSTM encoder hidden state size is 64. The linear embedding layers applied to all inputs of both encoders, and the first input of the decoder at $t = T_{obs}$, produce a 16-dimensional vector from the input coordinates. The linear embedding layer at the decoder's output produces a vector of $6 \times K$, where K is the number of components in the GMM, set as 6 for all experiments. Both MLPs have a hidden layer of size 64 and use ReLU activation. The network is trained initially for 10 epochs using only the negative log-likelihood loss L_{lh} , before training adversarially using both loss functions for a further 90 epochs. This initial training is implemented in order to encourage the Generator to produce sensible results before comparison to the ground truth by the Discriminator, and also allows training to converge in significantly fewer iterations. All training is performed using Adam optimiser with a batch size of 32 and initial learning rate of 0.001. The α weighting applied to L_{lh} in Eq. 8 is chosen as 0.1.

IV. EXPERIMENTS

We conduct two experiments in order to validate our method's effectiveness both with and without a vehicle feature input. Firstly, we evaluate our model without any vehicle feature input on two publicly available datasets of real world interacting pedestrian crowds, ETH [26] and UCY [27]. Next, we verify our model using a vehicle feature input on two datasets of interacting pedestrian crowds and vehicles. These include the publicly available dataset, Vehicle-Crowd Interaction DUT dataset (VCI) [28], and the USyd Campus Dataset (USyd) [1].

Metric	Dataset	Lin	CVM	SGAN	SRLSTM	Ours	
						PSGAN	PCGAN
ADE	ETH-Univ	0.50 / 0.79	0.48 / 0.70	0.51 / 0.81	0.43 / 0.65	0.45 / 0.68	0.43 / 0.65
	ETH-Hotel	0.35 / 0.39	0.28 / 0.33	0.55 / 0.67	0.24 / 0.42	0.52 / 0.64	0.59 / 0.64
	UCY-Univ	0.56 / 0.82	0.34 / 0.56	0.56 / 0.78	0.38 / 0.53	0.34 / 0.55	0.49 / 0.57
	UCY-Zara01	0.41 / 0.62	0.28 / 0.46	0.46 / 0.63	0.28 / 0.43	0.25 / 0.43	0.25 / 0.40
	UCY-Zara02	0.53 / 0.77	0.23 / 0.35	0.35 / 0.56	0.24 / 0.32	0.27 / 0.37	0.22 / 0.34
FDE	ETH-Univ	0.88 / 1.57	0.87 / 1.34	0.95 / 1.72	0.80 / 1.26	0.84 / 1.34	0.81 / 1.25
	ETH-Hotel	0.60 / 0.72	0.40 / 0.62	0.49 / 1.71	0.45 / 0.90	1.13 / 1.45	1.10 / 1.40
	UCY-Univ	1.01 / 1.59	0.71 / 1.20	1.20 / 1.70	0.81 / 1.17	0.71 / 1.23	0.89 / 1.24
	UCY-Zara01	0.74 / 1.21	0.57 / 0.99	0.99 / 1.38	0.60 / 0.93	0.53 / 0.87	0.53 / 0.89
	UCY-Zara02	0.95 / 1.48	0.47 / 0.75	0.75 / 1.21	0.51 / 0.73	0.56 / 0.76	0.45 / 0.77
MHD	ETH-Univ	0.48 / 0.66	0.40 / 0.57	0.44 / 0.66	0.38 / 0.54	0.40 / 0.59	0.38 / 0.55
	ETH-Hotel	0.33 / 0.33	0.20 / 0.27	0.22 / 0.69	0.22 / 0.37	0.45 / 0.56	0.51 / 0.55
	UCY-Univ	0.52 / 0.76	0.31 / 0.48	0.48 / 0.67	0.34 / 0.45	0.30 / 0.49	0.41 / 0.50
	UCY-Zara01	0.39 / 0.55	0.24 / 0.40	0.40 / 0.52	0.26 / 0.36	0.23 / 0.37	0.23 / 0.35
	UCY-Zara02	0.47 / 0.71	0.23 / 0.31	0.31 / 0.49	0.22 / 0.31	0.25 / 0.33	0.20 / 0.31

TABLE I: Quantitative results of tested methods on all non-vehicle datasets. For each dataset, we compare results across two prediction lengths of 8 and 12 timesteps (3.2 and 4.8 secs), showing Average Displacement Error (ADE), Final Displacement Error (FDE), and Modified Hausdorff Distance (MHD) in meters.

Metrics	Dataset	Lin	CVM	SGAN	SRLSTM	Ours	
						PSGAN	PCGAN
ADE	USyd	0.16	0.13	0.16	0.11	0.11	0.11
	VCI	0.11	0.09	0.12	0.08	0.12	0.08
FDE	USyd	0.30	0.24	0.31	0.22	0.21	0.21
	VCI	0.23	0.18	0.22	0.16	0.20	0.15
MHD	USyd	0.12	0.09	0.12	0.09	0.08	0.09
	VCI	0.09	0.07	0.09	0.07	0.09	0.07

TABLE II: Quantitative results of tested methods on both vehicle datasets. We compare results using a prediction length of 12 timesteps (1.0 second (VCI) and 1.2 seconds (USyd)). ADE, FDE and MHD are shown in meters.

A. Datasets

ETH and UCY contain 5 crowd scenes: ETH-Univ, ETH-Hotel, UCY-Zara01, UCY-Zara02, and UCY-Hotel. Each dataset is converted to world coordinates with an observation frequency of 2.5 Hz, similar to [2]. We deal with the ETH-Univ frame rate issue addressed by [24] similarly by treating every 6 frames as 0.4s rather than 10 frames, and retrain all comparative models for this scene.

USyd is collected on a weekly basis by [1] from March 2018 over the University of Sydney campus and surroundings. The dataset contains over 52 weeks drives and covers various environmental conditions. Since our research work primarily focuses on predicting socially plausible future trajectories of pedestrians under the influence of one vehicle, we select from the dataset 17 scenarios in an open large area with high pedestrian activity. Pedestrians are detected by fusing YOLO [29] classification results and LiDAR point clouds from vehicle onboard sensors, as illustrated in Fig. 1. The GMPHD [30] tracker is used to automatically label the trajectories of pedestrians. To increase the diversity of data available for training models, we apply data augmentation by flipping 2D coordinates randomly. Due to limitations regarding the length of time agents are observed in this dataset, we use an observation frequency of 10 Hz, rather than downsampling to be comparable to Experiment 1.

VCI proposed by [28], contains two scenes of labelled video from a birds eye view of vehicle-crowd interactions, recorded at 24 Hz. We downsample this dataset to 12 Hz in order to make results comparable with the USyd dataset. We remove sequences which contain more than one vehicle.

B. Evaluation Metrics and Baselines

1) *Metrics*: Similar to prior work [9] we included two error metrics: Average Displacement Error (ADE) and Final Displacement Error (FDE). However, as discussed by Zyner et al. [15], these commonly used measures do not consider outliers throughout the prediction and penalize misalignment in time and space equally. This can result in a prediction with an incorrect speed profile but correct direction having a similar error as a prediction with the completely wrong direction, which is a significantly worse result. As such, Modified Hausdorff Distance (MHD) [31], which does not suffer this issue, is also included as an evaluation metric.

The metrics used are as follows:

- ADE: Average Euclidean distance between ground truth and prediction trajectories over all predicted time steps.
- FDE: Euclidean distance between ground truth and prediction trajectories for the final predicted time step.
- MHD: A measure of similarity between trajectories, determining the largest distance from each predicted point to any point on the ground truth trajectory.

2) *Baseline Comparisons*: We compare our model against the following baseline and state of the art methods:

- Lin: A linear regression of pedestrian motion over each dimension.
- CVM: Constant Velocity Model proposed by [16].
- Social GAN (SGAN) [2]: LSTM encoder-decoder with a social pooling layer, trained as a GAN.
- SR-LSTM [24]: LSTM based model using a State Refinement module.

Additionally, we perform an ablation study of our method, comparing our model using a social pooling layer as pro-

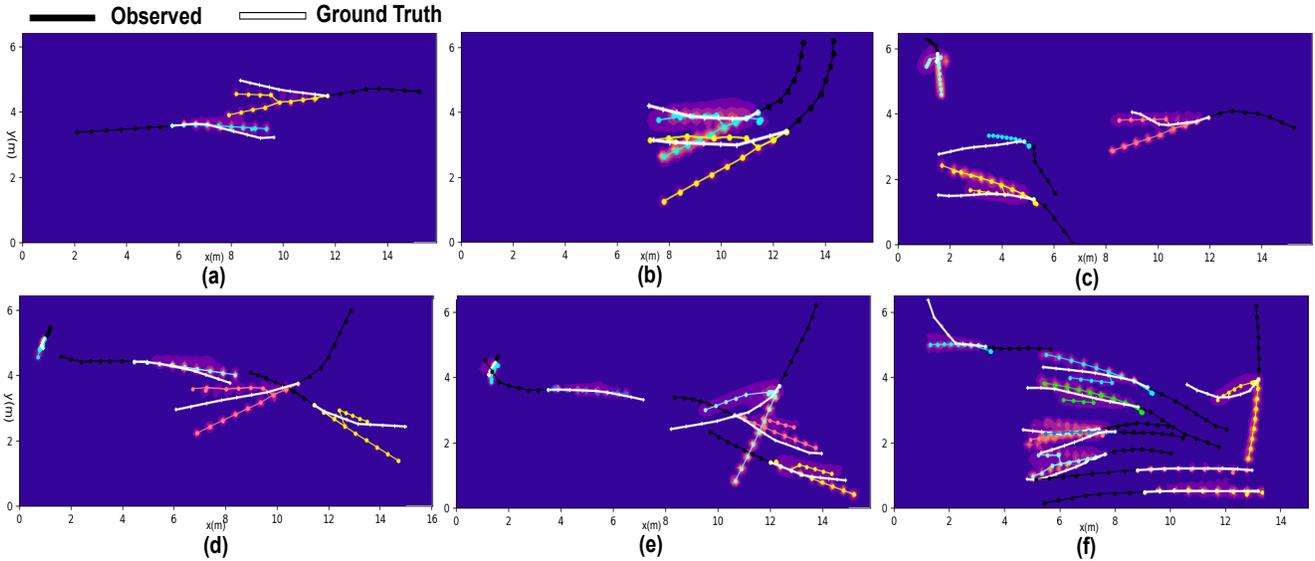


Fig. 4: The predicted modal path trees of MultiPAC are shown in a different colour for each pedestrian, over the probabilistic output of the Generator. Example interactions are from the ETH and UCY datasets, using PCGAN trained without vehicle feature input. Multimodal output is clear in examples in which pedestrians may take one of multiple possible future paths to avoid the collision. Example (a) displays two likely paths that the yellow agent might have taken as the pedestrians approach each other. Example (b) similarly shows multimodal possibilities, including the pedestrians continuing to turn, or to start travelling forwards. Examples (c) through (f) demonstrate similar behaviour in larger pedestrian crowds.

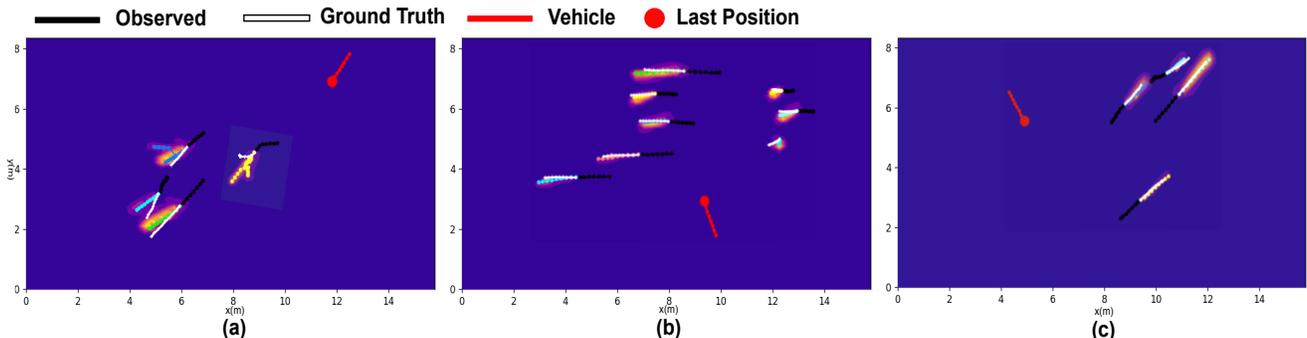


Fig. 5: Predicted pedestrian trajectories using PCGAN trained with vehicle feature input on the VCI dataset. Example (a) illustrates a scene of a vehicle approaching pedestrians from behind, displaying expected multimodal reactions of the pedestrians to either continue forwards at increased speed or move aside. Examples (b) and (c) further illustrate this concept, showing how the direction of the vehicle approach can impact the pedestrians’ reaction.

posed in [2] (PSGAN), and the model trained instead using our GVA module for social pooling (PCGAN).

As SGAN requires a random noise input for generation, we sample this method 10 times, returning the average error of all samples, as opposed to [2], where the sample with the best error compared to the ground truth was used.

C. Methodology

For all evaluations on our probabilistic methods, PSGAN and PCGAN, we apply MultiPAC to the output of the Generator to find all modal paths, using the predicted path with the highest probability to compute the error.

Experiment 1: Similar to [9], we train on four datasets and evaluate on the remaining one. We observe the last 8 timesteps of each trajectory (3.2 seconds) and predict for the

next 8 (3.2 seconds) and 12 (4.8 seconds) timesteps.

Experiment 2: Each dataset is split into non-overlapping train, validation and test sets in ratios of 60%, 20% and 20%. We observe 8 timesteps of each trajectory (0.67 seconds (VCI) and 0.8 seconds (USyd)) and predict for the next 12 timesteps (1.0 second (VCI) and 1.2 seconds (USyd)).

V. RESULTS AND DISCUSSION

A. Quantitative Evaluation

Experiment 1: Table.1 compares results for all methods on the ETH and UCY datasets. Our adversarial approaches PSGAN and PCGAN clearly outperform the previous sampling-based adversarial approach [2] demonstrating that the use of a direct probabilistic generator output can

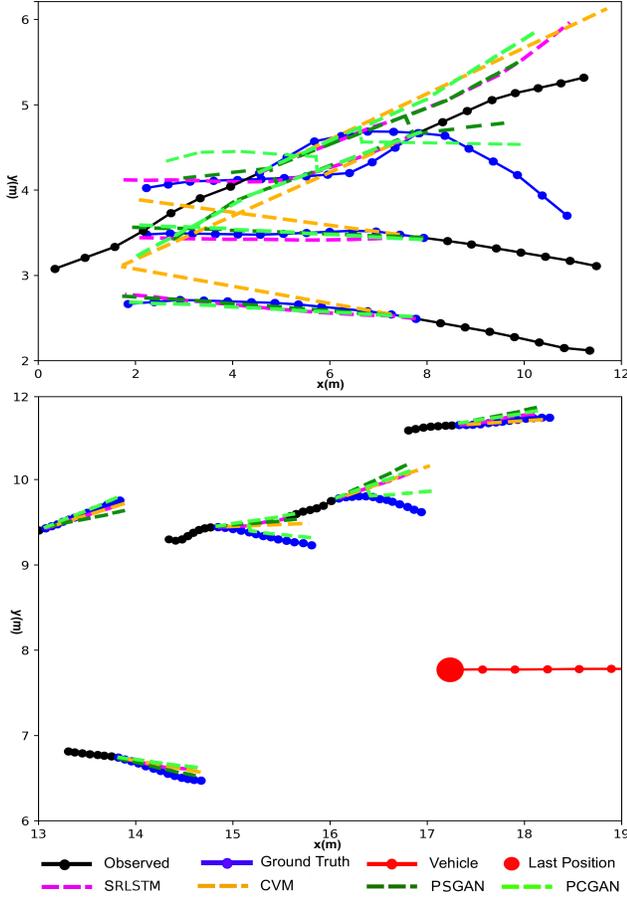


Fig. 6: Comparison of methods on UCY-Zara01 (top) and VCI (bottom) showing the entire modal path tree for both PSGAN and PCGAN. Whilst CVM and SRLSTM outperform our methods on some datasets, our multimodal output better represents uncertainty in crowd interactions, demonstrated in the top example where the possibility that oncoming pedestrians could avoid each other in two different ways is reflected in the branching modal path trees. The bottom illustrates how PCGAN improves predictions in the presence of a vehicle compared to PSGAN, accounting for the impact of the vehicle’s motion on pedestrians’ motion.

improve performance in the problem of trajectory prediction. Additionally, PCGAN and PSGAN achieve comparable or improved performance in 17 out of 30 metrics compared to prior methods suggesting that our probabilistic GAN approach can improve trajectory prediction performance in certain crowd interactions. Even when used without vehicle feature input we can see that the inclusion of the Gvat for social pooling in PCGAN improves the performance in the majority of tests compared to PSGAN. However, on both ETH-Hotel and UCY-Univ we find that PSGAN outperforms PCGAN. On these two datasets, CVM also performs well, suggesting that there may be fewer pedestrian interactions involved allowing more linear models to achieve improved results. Schöller et al. [16] demonstrated the effectiveness of CVM, and we find that this result still holds even when

limited to prediction periods of 8 and 12 timesteps. SGAN [2] performs poorly for all tested datasets when limited to using the average error over multiple samples, as opposed to using the best sample error compared to the ground truth. This result is similar to that obtained in [24], where SGAN was not found to perform well when limited to a single sample. Whilst this may be a result of SGAN sampling between multiple future modal paths, our method PSGAN, which extends SGAN for direct probabilistic output, demonstrates that by being able to estimate the likelihood of each modal path, we can greatly decrease the error of the adversarially trained method obtained for all metrics. Unlike both SGAN and our methods, State Refinement LSTM (SRLSTM) [24] pools across the pedestrian hidden states found from the most recent observation. Whilst only being comparable for predictions of 12 timesteps, this method performs well for all datasets, confirming the importance of using the most recently available information for predictions.

Experiment 2: Table II outlines the performance of all compared methods on the VCI and USyd datasets, both of which contain pedestrian-vehicle interactions. These results again highlight how using a probabilistic output during adversarial training can improve prediction results, with both of our methods, PSGAN and PCGAN, improving upon SGAN. Importantly, we can see that by including the vehicle feature input in Gvat pooling we can achieve significant improvements on the VCI dataset, with PCGAN significantly outperforming PSGAN, and outperforming or equalling SRLSTM on all metrics. CVM and PSGAN score well for the MHD metric, suggesting that these methods are likely incorrectly predicting the speed profile of pedestrian trajectories, but correctly predicting the direction.

B. Qualitative Evaluation

Experiment 1: Fig. 4 demonstrates realistic behaviours between pedestrians, producing results that reflect the actual probabilistic and multimodal nature of crowd interactions. (a) and (b) both reflect the ambiguity expected during an interaction between two pedestrians. The two possible trajectories that can be taken to avoid an oncoming pedestrian are clearly displayed in the modal paths of example (a), where one branch of the modal path tree matches the actual trajectory taken. This situation is again seen in Fig. 6 (top), where PSGAN (dark green) and PCGAN (light green) are able to accurately predict the turning of oncoming pedestrians with branching modal paths, whilst both SRLSTM (pink) and CVM (yellow) do not account for this ambiguity. Likewise, example (b) reflects the possibility that the two pedestrians might continue turning together, or instead travel forwards beside each other. Additional examples extend these ideas to more crowded scenes, with multiple pedestrians displaying the similar multimodal and uncertain interactions. In Fig. 4 (b), whilst there exists clear dependency between the two predicted forking modal path trees, our model does not currently have the ability to determine this relationship, and so cannot predict which branch an agent will take even with knowledge of the true path of a neighbouring agent.

Experiment 2: The extension of our approach to include a vehicle allows the modelling of interactions in shared pedestrian-vehicle environments, predicting crowd response in the presence of a vehicle as shown in *Fig.5*. Experiment 2 uses a shorter timestep, of only 0.1 second on the USyd, and 0.083 seconds for VCI. As is expected over this shorter time we don't see as significant interactions, reflected in near-linear ground truth in both *Fig.5* and *Fig.6* (bottom). However, we can still see clear multimodal predictions in certain interactions, including when the vehicle is approaching pedestrians from behind as in *Fig.6* (a), where the closest pedestrian responds by beginning to move to the side. This interaction is reflected in the predicted modal paths, although the sideways direction is predicted in the wrong direction. In *Fig.6* (bottom) we also see how only PCGAN accounts for the vehicle's influence on the pedestrians, correctly predicting the possibility that the pedestrians will return to their original motion once the vehicle has passed.

VI. CONCLUSION

Our work shows how a direct multimodal probabilistic output can be generated in an adversarial network for pedestrian trajectory prediction, outperforming existing methods including sampling-based approaches. We additionally show how the presence of an autonomous vehicle can be considered through the introduction of a novel GVAT pooling mechanism. By comparing our work to [2], a non-probabilistic GAN used for trajectory prediction, we have shown that our probabilistic approach clearly benefits adversarial training for this problem. Our work focuses on how a single vehicle can operate away from the lane-based structure of a road, examining crowd interactions to enable safer decisions, however could in future be extended for use with multiple vehicles through inclusion of all vehicles as nodes, removing the z term from *Eq. 9* and replacing W_r and W_u with a different set of weights for each agent type pair to learn relationship dynamics.

ACKNOWLEDGMENT

This work has been funded by the Australian Centre for Field Robotics(ACFR), University of Michigan / Ford Motors Company Contract "Next generation Vehicles", Transport for New South Wales (TfNSW) and iMOVE CRC and supported by the Cooperative Research Centres program, an Australian Government initiative.

REFERENCES

- [1] W Zhou, J S Berrio Perez, C De Alvis, M Shan, S Worrall, J Ward, and E Nebot. The USyd Campus dataset, 2019. 1, 4, 5
- [2] A Gupta, J Johnson, L Fei-Fei, S Savarese, and A Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *CVPR*, 2018. 1, 2, 5, 6, 7, 8
- [3] V Kosaraju, A Sadeghian, R Martín-Martín, I Reid, S H Rezatofighi, and S Savarese. Social-BiGAT: Multimodal Trajectory Forecasting using Bicycle-GAN and Graph Attention Networks. In *Conference on Neural Information Processing Systems*, 2019. 1, 2, 3
- [4] P Veličković, G Cucurull, A Casanova, A Romero, P Lio, and Y Bengio. Graph Attention Networks. *ICLR*, 2017. 1, 2, 3
- [5] N Lee, W Choi, P Vernaza, C B Choy, P HS Torr, and M Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *CVPR*, 2017. 2

- [6] S Yi, H Li, and X Wang. Understanding pedestrian behaviors from stationary crowd groups. In *CVPR*, 2015. 2
- [7] W Zhi, L Ott, and F Ramos. Kernel trajectory maps for multi-modal probabilistic motion prediction. *CoRL*, 2019. 2
- [8] D Helbing and P Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995. 2
- [9] A Alahi, K Goel, V Ramanathan, A Robicquet, L Fei-Fei, and S Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *CVPR*, 2016. 2, 5, 6
- [10] A Vemula, K Muelling, and J Oh. Social Attention : Modeling Attention in Human Crowds. In *ICRA*, 2018. 2
- [11] S S. Becker, R Hug, W Hubner, and M. Arens. RED: A simple but effective Baseline Predictor for the TrajNet Benchmark. *ECCV Workshop on Anticipating Human Behavior*, 2018. 2
- [12] T Salzmann, B Ivanovic, P Chakravarty, and M Pavone. Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control. *arXiv:2001.03093*, 2020. 2
- [13] Y Hu, W Zhan, L Sun, and M Tomizuka. Multi-modal probabilistic prediction of interactive behavior via an interpretable model. In *IEEE Intelligent Vehicles Symposium (IV)*, 2019. 2
- [14] X Huang, S McGill, J A DeCastro, B C Williams, L Fletcher, J J Leonard, and G Rosman. DiversityGAN: Diversity-aware vehicle motion prediction via latent semantic sampling. *arXiv:1911.12736*, 2019. 2
- [15] A Zyner, S Worrall, and E Nebot. Naturalistic driver intention and path prediction using recurrent neural networks. *IEEE Transactions on Intelligent Transportation Systems (ITSC)*, 2019. 2, 3, 4, 5
- [16] C Schöllner, V Aravantinos, F Lay, and A Knoll. What the Constant Velocity Model Can Teach Us About Pedestrian Motion Prediction. *ICRA*, 2020. 2, 5, 7
- [17] I Goodfellow, J Pouget-Abadie, M Mirza, B Xu, D Warde-Farley, S Ozair, A Courville, and Y Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*. 2014. 2
- [18] A Sadeghian, V Kosaraju, A Sadeghian, N Hirose, S. H Rezatofighi, and S Savarese. SoPhie: An Attentive GAN for Predicting Paths Compliant to Social and Physical Constraints. In *CVPR*, 2018. 2
- [19] B Ivanovic and M Pavone. The Trajectron: Probabilistic Multi-Agent Trajectory Modeling With Dynamic Spatiotemporal Graphs. In *ICCV*, 2019. 2
- [20] R Chandra, U Bhattacharya, A Bera, and D Manocha. TraPHic: Trajectory Prediction in Dense and Heterogeneous Traffic Using Weighted Interactions. In *CVPR*, 2019. 2
- [21] S Eiffert and S Sukkarieh. Predicting responses to a robot's future motion using generative recurrent neural networks. In *Australian Conference on Robotics and Automation (ACRA)*, 2019. 2
- [22] Y Ma, X Zhu, S Zhang, R Yang, W Wang, and D Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In *AAAI Conference on Artificial Intelligence*, 2019. 2
- [23] A Jain, A Zamir, S Savarese, and A Saxena. Structural-RNN: Deep learning on spatio-temporal graphs. In *CVPR*, 2016. 2
- [24] P Zhang, W Ouyang, P Zhang, J Xue, and N Zheng. SR-LSTM: State Refinement for LSTM towards Pedestrian Trajectory Prediction. In *CVPR*, 2019. 3, 5, 7
- [25] M Ester, H Krieger, J Sander, and X Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Int. Conference on Knowledge Discovery and Data Mining*, 1996. 4
- [26] S Pellegrini, A Ess, K Schindler, and L Van Gool. You'll never walk alone: Modeling social behavior for multi-target tracking. In *Int. Conference on Computer Vision (ICCV)*, 2009. 4
- [27] A Lerner, Y Chrysanthou, and D Lischinski. Crowds by example. In *Computer graphics forum*, volume 26, pages 655–664, 2007. 4
- [28] D Yang, L Li, K Redmill, and Ü Özgüner. Top-view trajectories: A pedestrian dataset of vehicle-crowd interaction from controlled experiments and crowded campus. *IEEE Intelligent Vehicles Symposium*, 2019. 4, 5
- [29] J Redmon, S Divvala, R Girshick, and A Farhadi. You only look once: Unified, real-time object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [30] B N Vo and W K Ma. The Gaussian mixture Probability Hypothesis Density Filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, 2006. 5
- [31] A Odena, C Olah, and J Shlens. Conditional image synthesis with auxiliary classifier GANs. In *ICML*, 2017. 5