

Euro-PVI: Pedestrian Vehicle Interactions in Dense Urban Centers

Apratim Bhattacharyya¹ Daniel Olmeda Reino² Mario Fritz³ Bernt Schiele¹

¹Max Planck Institute for Informatics, Saarland Informatics Campus

²Toyota Motor Europe

³CISPA Helmholtz Center for Information Security

Abstract

Accurate prediction of pedestrian and bicyclist paths is integral to the development of reliable autonomous vehicles in dense urban environments. The interactions between vehicle and pedestrian or bicyclist have a significant impact on the trajectories of traffic participants e.g. stopping or turning to avoid collisions. Although recent datasets and trajectory prediction approaches have fostered the development of autonomous vehicles yet the amount of vehicle-pedestrian (bicyclist) interactions modeled are sparse. In this work, we propose Euro-PVI, a dataset of pedestrian and bicyclist trajectories. In particular, our dataset caters more diverse and complex interactions in dense urban scenarios compared to the existing datasets. To address the challenges in predicting future trajectories with dense interactions, we develop a joint inference model that learns an expressive multi-modal shared latent space across agents in the urban scene. This enables our Joint- β -cVAE approach to better model the distribution of future trajectories. We achieve state of the art results on the nuScenes and Euro-PVI datasets demonstrating the importance of capturing interactions between ego-vehicle and pedestrians (bicyclists) for accurate predictions.

1. Introduction

Notwithstanding recent progress in the development of reliable self-driving vehicles, dense inner city environments remain challenging. One of the key components for the success of self-driving vehicles in dense urban environments is anticipation [4, 22]. Anticipating the motion of traffic participants in dense urban environments is made especially challenging due to the effect of interactions between different agents. For example, a pedestrian might turn onto the road to avoid an obstacle on the sidewalk which requires the vehicle to stop (Fig. 1c). Alternately, a pedestrian attempting to cross the road ahead of the ego-vehicle might continue or stop depending upon the distance and velocity of the vehicle

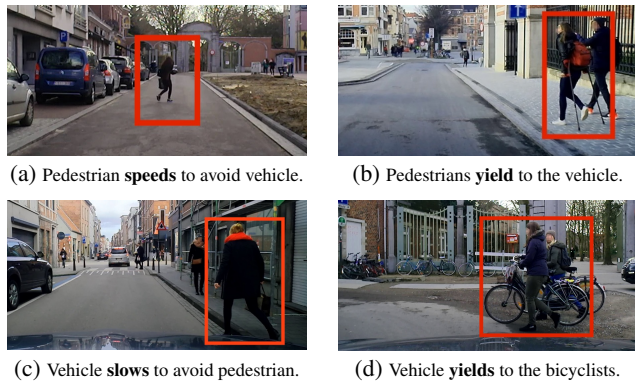


Figure 1. Examples of interactions between the ego-vehicle and pedestrians (bicyclists) in Euro-PVI.

(c.f. Fig. 2). Thus, interactions add significant complexity to the distribution of the likely future trajectories which is highly multi-modal.

Recently, datasets like nuScenes [6], Argoverse [8], or Lyft L5 [18] have greatly aided the development of trajectory prediction methods. However, these datasets are primarily focused on trajectories of vehicles and vehicle-vehicle interactions – collected to a large extent on multi-lane roads in North America or Asia, with sparse interactions between the ego-vehicle and pedestrians or bicyclists (e.g. Figure 4 in [18]). Therefore, they do not represent trajectories in dense urban environments well where interactions between the trajectories of agents are prominent. Such scenarios are particularly common in inner-city environments in Europe.

In this work, we propose the new European Pedestrian Vehicle Interaction (*Euro-PVI*) dataset¹ which is collected in a dense urban environment in Brussels and Leuven, Belgium. The Euro-PVI dataset contains a rich and diverse set of interactions between the ego-vehicle and pedestrians (bicyclists). Sequences are recorded near busy urban landmarks e.g. railway stations, narrow lanes or intersections (c.f. Figs. 1 to 3) where interactions are frequent and it is

¹The dataset is available at www.europvi.mpi-inf.mpg.de

therefore challenging to predict pedestrian (bicyclist) paths.

Further, in spite of the recent progress in trajectory prediction methods, accurately capturing the multi-modal distribution of future trajectories *e.g.* in dense urban environments remains challenging. Current state of the art [3, 28, 37] generative models for trajectory prediction encode interactions directly in the posterior. Thus, the latent space does not express interaction information from the input distribution which limits the accuracy of the generated future trajectories. To address this limitation, we develop a latent variable deep generative model which jointly models the distribution of future trajectories of the agents in the scene. Our *Joint- β -Conditional Variational Autoencoder (Joint- β -cVAE)* models a “shared” latent space between agents, to better capture the effect of interactions in the latent space and accurately represent the multi-modal distribution of trajectories.

Our contributions are, 1. We propose Euro-PVI, a novel dataset of pedestrian and bicyclist trajectories recorded in Europe with dense interactions with the ego-vehicle. 2. Our dataset facilitates research on dense interactions as we show that – in contrast to prior datasets – there is a pronounced performance gap between methods that model vehicle-pedestrian-interaction vs not. 3. To this end, we develop a latent generative model – Joint- β -cVAE – that models a shared latent space to better capture the effect of interactions on the multi-modal distribution of future trajectories. 4. Finally, we demonstrate state of the art performance on pedestrian (bicyclist) trajectory prediction on nuScenes and Euro-PVI.

2. Related Work

“On-board” Trajectory Prediction Datasets. Datasets like ETH/UCY [23] and Stanford Drone [32] are among the first datasets in the field of pedestrian trajectory prediction. However, they are recorded using a birds eye view camera or drone. In order to aid the development of autonomous driving capabilities, recent datasets have moved to a more realistic “on-board” setting – recorded from a (ego-)vehicle. The popular “on-board” datasets: nuScenes [6], Argoverse [8] and Lyft L5 [18] focus primarily on trajectories of nearby vehicles. nuScenes [6] and Argoverse [8] do not include annotated pedestrian (bicyclist) trajectories in their test set. Although Lyft L5 [18] includes pedestrian and bicyclist trajectories in the test set they are relatively rare (5.91% and 1.62% of all trajectories) as the chosen route does not include significant portions of dense urban environments. Moreover, in comparison to nuScenes [6] and Argoverse [8], Lyft L5 [18] has lower diversity in terms of locations as it recorded only along a fixed route (~ 6 km) in Palo Alto, California. Additionally, Lyft L5 [18] does not provide images from cameras and lidar point clouds, which are sources of rich contextual information. In contrast, PIE [33], TITAN [27] and TRAF [7] focus primarily on pedestrians (bicyclists). However, the trajectories are recorded as sequences of 2D

bounding boxes in the image plane and are not localized in the 3D world. The ApolloScapes [26] dataset does not include the trajectory of the ego-vehicle or contextual information *e.g.* images from cameras or lidar point clouds. Finally, note that these datasets are recorded either in North America or Asia and no large scale trajectory datasets are available for Europe. Euro-PVI is the first large scale dataset recorded in Europe dedicated to the task of trajectory prediction and unlike the existing datasets focuses on interactions between the ego-vehicle and pedestrian (bicyclist).

Methods for Trajectory Prediction. Approaches which recognize the importance of interactions for the task of trajectory prediction date back to the Social Forces approach [16]. More recently Social-LSTM [1] proposes a social pooling mechanism to capture interactions. The social pooling mechanism has been extended in [10] for efficiency. Alternately, TrafficPredict [26] uses an instance layer to embed a spatio-temporal graph of interactions. MATF [45] uses a tensor fusion scheme to capture interactions while retaining the spatial structure of the scene. TraPHic [7] proposes weighted interactions. Car-Net [36] uses an attention scheme to integrate visual cues with interactions. Joint prediction of trajectories and activities is performed in [24]. Social-STGCNN [30] models the prediction scene as a graph and uses graph convolutions to capture interactions. Spatio-temporal graph transformer networks are used in [43]. A planning informed method is developed in [39]. However, these works largely assume a uni-modal future or a future conditioned using a set of possible maneuvers and cannot fully capture multi-modal distributions *e.g.* trajectories in dense urban environments.

Therefore, recent work has focused on the development of generative approaches to capture multi-modal distributions of trajectories, primarily using conditional GAN [29], conditional Normalizing flows [11] or conditional VAE [38] based models. Social-GAN [15] proposes to use social pooling in a conditional GAN setup to capture the effect of social interactions in the distribution of future pedestrian trajectories. Sophie [35] uses an attention mechanism in a conditional GAN setup to capture the effect of interactions. Graph attention networks are employed by [21] in a conditional GAN setup. Normalizing flows are combined with a context attention mechanism in [31] to improve expressiveness in the latent space. HBA-Flow [5] improves trajectory prediction with a Haar-wavelet based decomposition. DESIRE [22] uses a RNN refinement module over the plain cVAE setup. Predictions that “personalize” to agent behaviour are proposed in [12]. The standard cVAE objective is improved in [4] for better match between the predicted and true distributions. A social graph network is used to condition a cVAE in [44]. It is shown in [28] that conditioning the cVAE additionally on the goal state can significantly improve accuracy. Finally, Trajectron++ [37] uses a pooled scene graph in a cVAE framework to capture interactions. The resulting

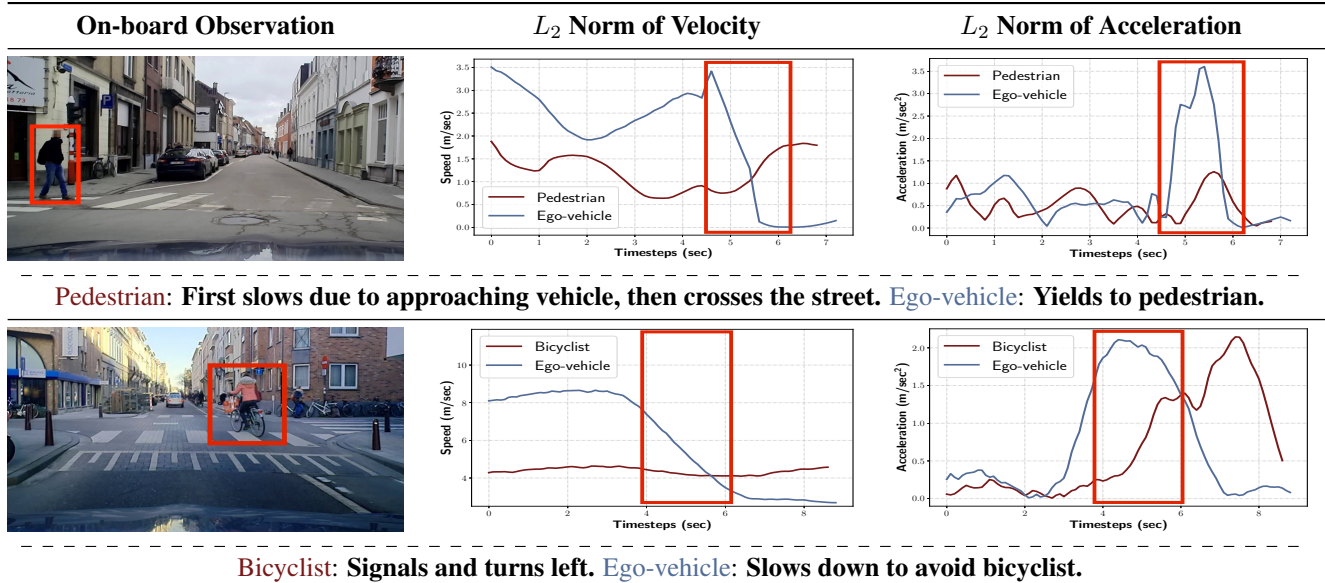


Figure 2. Examples of interactions in Euro-PVI. Spikes in the magnitude (L_2 norm) of acceleration resulting from interactions are marked.

approach leads to state of the art pedestrian trajectory prediction both in the ETH/UCY datasets and in the “on-board” setting of nuScenes. Although cVAE based methods achieve state of the art performance, the latent space in the above mentioned works does not encode the effect of interactions. This makes it challenging to accurately capture the trajectory distribution *e.g.* in dense urban environments where these interactions have a significant effect. In this work, we propose to use a shared latent space across agents to better capture the effect of interactions to generate accurate future trajectories.

3. The Euro-PVI Dataset

In this section, we introduce our Euro-PVI dataset to advance the task of “on-board” trajectory prediction especially in dense urban environments. The dataset focuses on the role of the ego-vehicle - pedestrian (bicyclist) interactions present in a scene to predict the future trajectories in dense urban environments. We first concretely define an “interaction”, which guides our data collection process and helps us select relevant sequences. Next, we provide details of the sensor setup and the data collection process. We then compare Euro-PVI to the existing datasets with respect to the density of interactions and provide detailed dataset statistics.

Interactions. We define an interaction between the ego-vehicle and a traffic participant (*e.g.* pedestrian or bicyclist) as an event where the presence of either (or both) the ego-vehicle or the traffic participant causes a change in velocity (change of speed/direction of motion *i.e.* acceleration) of the other. Examples of interactions include, 1. The ego-vehicle yielding to a pedestrian at a crosswalk (Fig. 2 top).

2. Non-verbal communication causes the ego-vehicle to slow down, as to avoid a bicycle which wants to turn (Fig. 2 bottom). We aim to record sequences which contain dense interactions between the ego-vehicle and vulnerable traffic participants, in particular, pedestrians and bicyclists.

Drive Planning and Scene Selection. Euro-PVI is recorded in the dense urban scene of Brussels and Leuven, Belgium. The ego-vehicle is equipped with a 360° lidar, a positioning system and a set of front-facing synchronized cameras (see supplemental material). The ego-vehicle is crewed by a pilot and a co-pilot. The pilot is instructed to drive freely over a predetermined area and to re-visit locations at times when dense concentration of pedestrians (bicyclists) is to be expected, such as transport hubs during peak hour. The duration of the driving sessions were up to 8 hours per day, over the course of two weeks. The co-pilot is tasked with identifying interactions and tagging the event. In the case of changes in trajectory or velocity of the ego-vehicle, the co-pilot asks the pilot for confirmation. The pilot is also instructed to spontaneously indicate that an interaction has happened. In Fig. 3 we show the geographical distributions of the trajectories in two example locations. We see that there is a high density of trajectories located around busy urban landmarks *e.g.* the railway station. We also show that interactions are not confined to road intersections with crosswalks where pedestrian (bicyclist) trajectories are simpler to predict, but also occur at locations without crosswalks (*i.e.* without designated crossing areas for pedestrians/bicyclists) where trajectories are more challenging to predict.

Interactions in Urban Environments. We now compare Euro-PVI to existing datasets for trajectory prediction with




Figure 3. Examples of aggregated spatial distribution of trajectories of pedestrians and cyclists around intersections and urban landmarks.




Figure 4. Left: Cumulative distribution sorted by distance to the ego-vehicle. Close proximity of the ego-vehicle and pedestrians (bicyclists) in Euro-PVI indicate dense traffic scenarios where interactions are likely. Right: Maximum acceleration sorted by distance to the ego-vehicle. High acceleration in close proximity of the ego-vehicle and pedestrians (bicyclists) indicate high likelihood of interactions.

respect to the density of interactions, in particular to the two largest datasets – nuScenes [6] and Lyft L5 [18]. First, we compare the distances between the ego-vehicle and pedestrians (bicyclists) in the scene. Short distances are indicative of closely packed urban environments where interactions frequently occur. In Fig. 4 (left), we show the closest approach (proximity) of a pedestrian or bicyclist to the ego-vehicle. We see that in the nuScenes and Lyft L5 datasets, the majority of the pedestrians (bicyclists) do not approach the ego-vehicle closer than ~ 20 meters. Such large distances between the pedestrians (bicyclists) and the ego-vehicle, more than 4 typical car lengths, decreases the likelihood of interactions. In contrast, in Euro-PVI the majority of the pedestrians (bicyclists) approach the ego-vehicle is as close as ~ 8 meters. Such short distances between the ego-vehicle and pedestrians (bicyclists) are indicative of the densely packed urban environment in which Euro-PVI is recorded with increases the likelihood of interactions.

However, close proximity only increases the likelihood of interactions, but does not necessarily lead to interactions. Note that by definition, interactions lead to a change in velocity *i.e.* acceleration. Therefore, in Fig. 4 (right) we plot the maximum acceleration experienced by the pedestrian (bicyclist) and ego-vehicle with increasing distance to the ego-vehicle and a pedestrian (bicyclist) respectively. In case of nuScenes and Lyft L5 we do not see a strong dependence on distance. In contrast, in Euro-PVI both pedestrians

(bicyclists) and the ego-vehicle experience the maximum acceleration close to the point of closest approach. This is again strongly indicative of dense interactions in Euro-PVI.

Qualitative Examples of Interactions. We provide example interactions in Euro-PVI along with the resultant acceleration of the involved agents in Fig. 2, *e.g.* top row: the pedestrian first slows down due to the approaching ego-vehicle and at the same time, the ego-vehicle sees the pedestrian and yields. This is visible as a spike in the velocity and acceleration plots. Similar spikes in acceleration can be observed due to interactions in the other examples in Fig. 2. We provide additional examples in the supplemental material.

Additional Statistics. We report dataset statistics and available labels of Euro-PVI in Table 1. In addition to the annotated pedestrian (bicyclist) trajectories, Euro-PVI contains 83k camera images and the corresponding lidar point clouds along with synchronized IMU data. In terms of size Euro-PVI is competitive with nuScenes [6] and ApolloSpaces [26] and while Lyft L5 [18] is significantly larger, it does not provide labels *e.g.* camera images or lidar point clouds. Furthermore, Euro-PVI surpasses the the largest autonomous driving datasets collected in Europe *i.e.* CityScapes [9], KITTI [13], KITTI-360 [42] and A2D2 [14] in terms of number of instances of pedestrian (bicyclist) trajectories. CityScapes [9] and A2D2 [14], do not provide annotated 3D pedestrian or bicycle trajectories. KITTI [13] and KITTI-360 [42] contains mostly linear motion with sparse interactions and thus commonly not used for trajectory prediction [26, 34]. In fact, Euro-PVI is the first large scale dataset with dense interactions (Fig. 4) dedicated to trajectory prediction in Europe, to the best of our knowledge.

4. Joint- β -cVAE: A Joint Trajectory Model for Dense Urban Environments

Following the observations in Section 3, we find that vehicle-pedestrian (bicyclist) interactions are crucial to the task of future trajectory prediction in dense urban scenarios. In particular, inherent multi-modality of the distribution of future trajectories and the effect of interactions on this complex distribution, make accurately predicting the future trajectories in dense urban environments challenging.

Specifically, given a scene with n agents *e.g.* vehicles, pedestrians or bicyclists in a dense urban environment and the past observations $\mathbf{x}_i \in \mathbf{X}$ for each agent i , we model the future trajectories $\mathbf{y}_i \in \mathbf{Y}$ for each agent $i \in \{1, \dots, n\}$ in the scene. Here the past observations \mathbf{x}_i include the past trajectories and the past context corresponding to the past trajectory sequence. Prior work [3, 4, 28, 37] models the conditional distribution $p_\theta(\mathbf{y}_i|\mathbf{X})$ parameterized by θ of the future trajectories \mathbf{y}_i using the latent variables \mathbf{z}_i in the

Dataset	Location	Scenes	Length (hrs)	Trajectory Instances		Labels			
				Pedestrians	Bicyclists	Front Camera	Lidar	Seg. Maps	IMU
nuScenes [6]	North Am. & Asia	1000	5.5	9142	550	✓	✓	✓	✓
ApolloScapes [26]	Asia	54	1.7	3065	1827*	x	x	x	x
Lyft L5 [18]	North Am.	170×10^3	1118	605×10^3	77×10^3	x	x	x	✓
CityScapes [9]	Europe	5000	2.5	0	0	✓	x	✓	✓
KITTI [13]	Europe	71	1.5	380	150	✓	✓	✓	✓
KITTI-360 [42]	Europe	9	2.2	262	82	✓	✓	✓	✓
A2D2 [14]	Europe	23	0.9	260	248	✓	✓	✓	✓
Euro-PVI (Ours)	Europe	1077	2.2	6177	1581	✓	✓	✓	✓

Table 1. Comparison of dataset statistics. (Seg. Maps – semantic and/or instance segmentation, * includes motorbikes.)

standard conditional VAE formulation [17, 20, 38],

$$p_\theta(\mathbf{y}_i|\mathbf{X}) = \int p_\theta(\mathbf{y}_i|\mathbf{z}_i, \mathbf{X}) p_\theta(\mathbf{z}_i|\mathbf{x}_i) d\mathbf{z}_i. \quad (1)$$

Here, the distribution $p_\theta(\mathbf{z}_i|\mathbf{x}_i)$ assumes conditional independence of the latent variables $\{\mathbf{z}_1, \dots, \mathbf{z}_n\} \in \mathbf{Z}$ given the past observation \mathbf{x}_i of each agent. This assumption essentially ignores the motion patterns of interacting agents *i.e.* the ego-vehicle and other pedestrians (bicyclists) in the scene, which in real world dense urban scenarios is critical for the accurate prediction of future trajectories. The formulation in Eq. (1) therefore limits the amount of interactions between the agents that can be encoded in the latent space. Since, the latent variables \mathbf{z}_i are crucial for capturing diverse futures in such conditional models, it is important to express the effect of interactions in the latent space.

We now introduce our Joint- β -cVAE approach which aims to accurately capture the effect of interactions in the latent space for trajectory prediction in dense urban environments. Our proposed Joint- β -cVAE model, in contrast to prior conditional VAE based models [3, 4, 28, 37] encodes the joint distribution of the latent variables across all agents in the scene. This allows our Joint- β -cVAE model to encode the dependence of the future trajectory distribution on interacting agents in the latent space, leading to more accurate modeling of the multi-modal future trajectory distribution.

Formulation. Our Joint- β -cVAE in Fig. 5 models the joint distribution of future trajectories \mathbf{Y} across all n agents, using the latent variables $\mathbf{z}_i \in \mathbf{Z}$,

$$\begin{aligned} p_\theta(\mathbf{Y}|\mathbf{X}) &= \int p_\theta(\mathbf{Y}, \mathbf{Z}|\mathbf{X}) d\mathbf{Z} \\ &= \int \prod_i^n p_\theta(\mathbf{y}_i, \mathbf{z}_i|\mathbf{Z}_{<i}, \mathbf{Y}_{<i}, \mathbf{X}) d\mathbf{Z} \\ &= \int \prod_i^n p_\theta(\mathbf{y}_i|\mathbf{Z}_{\leq i}, \mathbf{Y}_{<i}, \mathbf{X}) p_\theta(\mathbf{z}_i|\mathbf{Z}_{<i}, \mathbf{Y}_{<i}, \mathbf{X}) d\mathbf{Z} \end{aligned} \quad (2)$$

where $p_\theta(\mathbf{Y}, \mathbf{Z}|\mathbf{X})$ is the joint distribution of the future trajectories and the latent variables for all agents. In the second step, without loss of generality, we auto-regressively

factorize the joint distribution over the n agents, where $\mathbf{Z}_{<i}, \mathbf{Y}_{<i}$ denotes the latent variables and trajectories for agents $\{1, \dots, i-1\}$. Note that, the factorization is agnostic to the choice of ordering of the agents. In contrast to Eq. (1), the prior distribution of the latent variables in Eq. (2) exhibits joint modeling of the latent variables \mathbf{z}_i , *i.e.* $p_\theta(\mathbf{z}_i|\mathbf{Z}_{<i}, \mathbf{Y}_{<i}, \mathbf{X})$.

We maximize the log-likelihood of the data under the model in Eq. (2) with variational inference using a joint variational posterior distribution $q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$.

The Joint Posterior. To encode rich latent spaces shared between the n agents which capture the effect of interactions in dense urban environments, we propose a joint posterior over all n agents in the scene, which auto-regressively factorizes,

$$q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{Y}) = q_\phi(\mathbf{z}_1|\mathbf{X}, \mathbf{Y}) \cdots q_\phi(\mathbf{z}_n|\mathbf{Z}_{<n}, \mathbf{X}, \mathbf{Y}). \quad (3)$$

The conditional distributions $q_\phi(\mathbf{z}_i|\mathbf{Z}_{<i}, \mathbf{X}, \mathbf{Y})$ corresponding to agent i are normal distributions whose mean and variances are functions of $\mathbf{Z}_{<i}, \mathbf{X}$ and \mathbf{Y} . Intuitively, given the past trajectories of the agents $\{1, \dots, n\}$ and the joint latent posterior distribution over the interacting agents $\{1, \dots, i\}$, the latent posterior distribution corresponding to agent i encoding its interactions can be inferred conditioned on the latent information of the interacting agents $\{1, \dots, i-1\}$. We use an attention mechanism inspired by [2], to model the dependence of the distribution of \mathbf{z}_i on $\mathbf{Z}_{<i}$ and $\{\mathbf{X}, \mathbf{Y}\}$. The attention weights on \mathbf{z}_j and $\{\mathbf{x}_j, \mathbf{y}_j\}$, for $j \neq i$, is additionally conditioned on the past observation and location of the agents (Fig. 5) – which allows to attend to agents j interacting with agent i .

In contrast, prior work [3, 4, 28, 37] employ conditionally independent posteriors $q_\phi(\mathbf{z}_i|\mathbf{x}_i, \mathbf{y}_i)$ across agents in a scene – encoding limited interactions in the latent space.

The Joint Prior. The prior term in Eq. (2), $p_\theta(\mathbf{z}_i|\mathbf{Z}_{<i}, \mathbf{Y}_{<i}, \mathbf{X})$, encodes the effect of interactions on the latent space of agent i through the dependence on $\mathbf{Z}_{<i}, \mathbf{X}$. In practice, we find that a simpler joint prior,

$$p_\theta(\mathbf{Z}|\mathbf{X}) = p_\theta(\mathbf{z}_1|\mathbf{X}) \cdots p_\theta(\mathbf{z}_n|\mathbf{Z}_{<n}, \mathbf{X}) \quad (4)$$

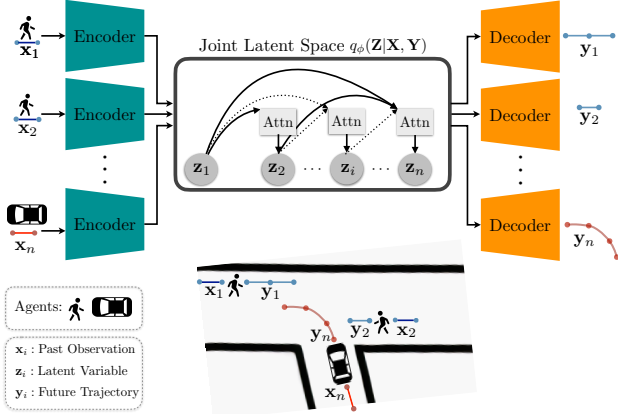


Figure 5. Our Joint- β -cVAE, which models a Joint latent space across all agents $\{1, \dots, n\}$ in the scene. The posterior latent distribution $q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$ factorizes auto-regressively and models the dependence of \mathbf{z}_i on $\{\mathbf{Z}_{<i}, \mathbf{X}, \mathbf{Y}\}$ using an attention mechanism.

is sufficient for rich latent spaces that captures interactions. We parameterize the prior as a conditional normal distribution, where the mean and variance depends on $\{\mathbf{Z}_{<i}, \mathbf{X}\}$.

The ELBO. As the standard log evidence lower bound (ELBO) for cVAEs, proposed in [20, 38], experiences several issues *e.g.* posterior collapse, we employ the ELBO formulation of β -VAE [17] to improve the representational capacity of the latent space and more accurately capture the effect of interactions in dense urban environments. With the formulation of the factorized variational distribution $q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$ in Eq. (3) and the joint prior distribution in Eq. (4), the ELBO is given by (details in the supplemental),

$$\log(p_\theta(\mathbf{Y}|\mathbf{X})) \geq \sum_i \mathbb{E}_{q_\phi} \log(p_\theta(\mathbf{y}_i|\mathbf{z}_{\leq i}, \mathbf{Y}_{<i}, \mathbf{X})) - \beta \sum_i D_{\text{KL}}(q_\phi(\mathbf{z}_i|\mathbf{Z}_{<i}, \mathbf{X}, \mathbf{Y}) || p_\theta(\mathbf{z}_i|\mathbf{Z}_{<i}, \mathbf{X})). \quad (5)$$

Additionally, we find it beneficial to model the observation noise σ^2 in the posterior distribution over the trajectories $p_\theta(\mathbf{y}_i|\mathbf{z}_{\leq i}, \mathbf{Y}_{<i}, \mathbf{X})$ as recommended in [25]. During training, we alternately optimize both the posterior and the prior distributions such that the ELBO is maximized [3, 41].

We find in practice that is it sufficient to condition the decoder $p_\theta(\mathbf{y}_i|\mathbf{z}_{\leq i}, \mathbf{Y}_{<i}, \mathbf{X})$ only on $\{\mathbf{Z}_{\leq i}, \mathbf{X}\}$. The rich latent space of our Joint- β -cVAE already models the effect of interactions, making it unnecessary to additionally condition on $\mathbf{Y}_{<i}$ for good performance.

5. Experiments

In this section, we 1. Demonstrate the effectiveness of our Joint- β -cVAE method. 2. Provide additional experimental evidence to better highlight the differences in the density of interactions between the ego-vehicle and pedestrians(bicyclists) Euro-PVI and current datasets. In order

to address the above points, in addition to Euro-PVI, we evaluate on nuScenes [6]. We choose nuScenes [6] as it is significantly larger compared to datasets like ApolloScapes [26] and more diverse in comparison to Lyft L5 [18], while possessing similar proximity/acceleration statistics (Fig. 4). We first evaluate our approach on nuScenes dataset followed by the evaluation on our proposed Euro-PVI dataset.

Evaluation Metrics. Following [37], we report, 1. Best of N (FDE): The final (euclidean) displacement error in meters using the best of $N = 20$ samples [15, 35, 37]. 2. KDE NLL: The (mean) negative log-likelihood of the groundtruth trajectory under the predicted distribution estimated using a Gaussian kernel [19, 40], computed using the code provided by [37]. Both these metrics aim to measure the match of the predicted trajectory distribution to the groundtruth distribution [3, 15, 22]. We provide results with the average (euclidean) displacement error in the supplemental material.

Implementation Details. We provide additional implementation details *e.g.* number of layers, hyper-parameters and code in the supplemental material.

5.1. nuScenes

Following [37], we split the training set into the training and validation splits. The original validation split is used as test set. We provide 1 - (*upto*) 5 secs of observation (historical context) and predict 3 seconds ahead [37]. We compare our Joint- β -cVAE approach to the following state of the art models: Social-GAN [15], Sophie [35] and Trajectron++ [37]. Additionally, in order to measure the density and influence of interactions between the ego-vehicle and pedestrians (bicyclists) on the trajectories in nuScenes (in comparison to Euro-PVI), we also evaluate the above methods without modeling ego-vehicle - pedestrians (bicyclists) interactions. Any significant difference in performance of these models would indicate the presence of dense ego-vehicle - pedestrian (bicyclist) interactions.

To illustrate the effectiveness of our Joint- β -cVAE, we also include two ablations of our Joint- β -cVAE model, 1. A simple cVAE model and, 2. A β -cVAE model, (neither of which can model interactions). These ablations are designed to show the effectiveness of our Joint- β -cVAE model in capturing interactions in the latent space.

We report results using both the Best of N and KDE NLL metrics in Table 2. The P-P and P-V columns in Table 2 indicate whether the pedestrian (bicyclist) - pedestrian (bicyclist) and pedestrian (bicyclist) - ego-vehicle interactions are modeled – using social pooling in case of Social-GAN [15], the attention mechanism in case of Sophie [35], the scene graph in case of Trajectron++ [37] and with a shared latent space in case of our Joint- β -cVAE model. Trajectron++ outperforms both the conditional GAN based Social-GAN and Sophie models – partly due to the better modeling of interaction with the scene graph compared to Social-GAN and

Method	Interactions		Best of $N = 20$ ↓			KDE NLL ↓		
	P-P	P-V	$t + 1$ sec	$t + 2$ sec	$t + 3$ sec	$t + 1$ sec	$t + 2$ sec	$t + 3$ sec
Social-GAN [15]	✓	–	0.04	0.11	0.21	-2.78	-1.40	-0.46
Social-GAN [15]	✓	✓	0.04	0.11	0.21	-2.80	-1.41	-0.48
Sophie [35]	✓	–	0.04	0.11	0.21	-2.59	-1.26	-0.41
Sophie [35]	✓	✓	0.04	0.11	0.21	-2.63	-1.27	-0.42
Trajectron++ [37]	✓	–	0.01	0.08	0.15	-5.55	-3.87	-2.69
Trajectron++ [37]	✓	✓	0.01	0.08	0.15	-5.58	-3.96	-2.77
cVAE	–	–	0.05	0.12	0.23	-2.51	-1.20	-0.21
β -cVAE [17]	–	–	0.01	0.08	0.17	-6.90	-4.10	-2.41
Joint- β -cVAE (Ours)	✓	–	0.01	0.06	0.13	-7.50	-4.53	-2.95
Joint- β -cVAE (Ours)	✓	✓	0.01	0.06	0.13	-7.55	-4.59	-2.98

Table 2. Evaluation on nuScenes. P-P and P-V: whether pedestrian - pedestrian or pedestrian - ego-vehicle interactions are modeled.

Method	Interactions		Best of $N = 20$ ↓			KDE NLL ↓		
	P-P	P-V	$t + 1$ sec	$t + 2$ sec	$t + 3$ sec	$t + 1$ sec	$t + 2$ sec	$t + 3$ sec
Social-GAN [15]	✓	–	0.14	0.36	0.65	-0.38	0.74	1.55
Social-GAN [15]	✓	✓	0.14	0.35	0.64	-0.41	0.73	1.52
Sophie [35]	✓	–	0.11	0.30	0.58	-1.53	-0.22	0.53
Sophie [35]	✓	✓	0.11	0.29	0.56	-1.71	-0.31	0.40
Trajectron++ [37]	✓	–	0.09	0.29	0.56	-2.75	-0.91	0.23
Trajectron++ [37]	✓	✓	0.09	0.28	0.54	-2.81	-1.00	0.15
cVAE	–	–	0.12	0.32	0.60	-1.40	-0.08	0.78
β -cVAE [17]	–	–	0.10	0.30	0.56	-2.61	-0.78	0.31
Joint- β -cVAE (Ours)	✓	–	0.09	0.29	0.53	-3.69	-1.29	0.02
Joint- β -cVAE (Ours)	✓	✓	0.09	0.27	0.51	-3.75	-1.38	-0.05
Joint- β -cVAE + {Camera, Lidar} (Ours)	✓	✓	<i>0.09</i>	<i>0.27</i>	<i>0.50</i>	-3.78	-1.41	-0.13

Table 3. Evaluation on Euro-PVI. P-P and P-V: whether pedestrian - pedestrian or pedestrian - ego-vehicle interactions are modeled.

Sophie. Also, note that Trajectron++ is built using a cVAE backbone. Thus, the performance advantage of Trajectron++ also illustrates the effectiveness of cVAE based models in capturing the distribution of future trajectories. We see that our Joint- β -cVAE outperforms Trajectron++. Additionally, our Joint- β -cVAE outperforms both the simple cVAE and β -cVAE ablations, illustrating that our Joint- β -cVAE model can effectively model interactions in the latent space. The performance advantage of our Joint- β -cVAE model over Trajectron++ shows the advantage of a joint latent space that can model the effect of interactions, in comparison to independent latent spaces which model the effect of interactions only as an additional condition to the decoder. Finally, across all models, we see that models which additionally encode pedestrian (bicyclist) - ego-vehicle interactions in nuScenes do not show a significant improvement in performance. This further lends support to the fact that pedestrian (bicyclist) - ego-vehicle interactions are sparse in the nuScenes dataset.

5.2. Euro-PVI

We now evaluate the different models for trajectory prediction on our novel Euro-PVI dataset. We use 792 sequences for training, 100 sequences for validation and 185 sequences for testing. The train/val/test splits do not share

pedestrian (bicyclist) instances. As on nuScenes, we compare our Joint- β -cVAE model with the Social-GAN [15], Sophie [35], Trajectron++ [37] models. We also include the cVAE and β -cVAE ablations (which cannot model interactions), to establish whether our Joint- β -cVAE approach can model interactions in the latent space. We follow a similar evaluation protocol as in nuScenes, where we predict trajectories up to 3 seconds into the future. However, we provide a shorter observation of 0.5 seconds as quick reactions are essential in dense traffic scenarios.

We report the results in Table 3. As on nuScenes, we observe that our Joint- β -cVAE approach outperforms the competing methods. The performance advantage over Trajectron++ again illustrates the advantage of the joint latent space over all agents in the scene versus an independent latent space which cannot model the effect of interactions. Our Joint- β -cVAE outperforms the cVAE and β -cVAE baselines, which illustrates that our Joint- β -cVAE model can model the effect of interactions successfully in the latent space. Additionally, the performance gain on Euro-PVI (0.03m, Best of $N = 20$) of our Joint- β -cVAE model over Trajectron++ is larger than in nuScenes. This shows that our Joint- β -cVAE model can better capture the complex distribution of pedestrians (bicyclists) trajectories in dense urban environments




Figure 6. Qualitative examples on Euro-PVI. We compare the Best of $N=20$ samples for Trajectron++ (red) and our Joint- β -cVAE (blue).

under the effect of interactions. We further show that performance can be improved by conditioning our Joint- β -cVAE model on visual features from the camera and lidar (also see supplemental material). We provide qualitative examples comparing our Joint- β -cVAE model to Trajectron++ in Fig. 6. We see in Fig. 6 (left) that Trajectron++ incorrectly predicts that the pedestrian will step onto the road, while our Joint- β -cVAE correctly predicts that due to the oncoming ego-vehicle the pedestrian avoids stepping onto the road. Similarly, in Fig. 6 (middle) our Joint- β -cVAE correctly predicts that the pedestrian quickly crosses the street due to the oncoming ego-vehicle and in Fig. 6 (right) the bicyclist merges in front of the ego-vehicle which slows down.

Finally, across all methods, we see the gain in performance (using both the Best of N and KDE NLL metrics) across methods when pedestrian (bicyclist) - ego-vehicle (P-V) interactions are modeled in Table 3 is larger than in nuScenes. This provides further evidence of dense pedestrian(bicyclist) - ego-vehicle interactions in Euro-PVI compared to the sparse interactions in nuScenes. Additionally, in Fig. 7 we plot the Best of N error of our Joint- β -cVAE model along with the β -cVAE ablation versus the distance of the trajectory from the ego-vehicle for both nuScenes and Euro-PVI. We see that in case of nuScenes, the error does not depend strongly on distance. This mirrors the results in Fig. 4 (right) which again suggests that the interactions between pedestrians and the ego-vehicle are sparse. In contrast, in case of our Euro-PVI dataset, the error is largest when the distance between the pedestrian (bicyclist) and the ego-vehicle is smallest *i.e.* at close encounters. This again suggests the presence of dense pedestrian (bicyclist) - ego-vehicle interactions in Euro-PVI.

Method	Best of $N=20 \downarrow$		
	$t+1$ sec	$t+2$ sec	$t+3$ sec
Trajectron++ [37]	0.10	0.35	0.63
Joint- β -cVAE (Ours)	0.10	0.33	0.61

Table 4. Transferring models trained on nuScenes to Euro-PVI.

Transferring Models from nuScenes. Finally, we experiment with transferring the best performing models on nuScenes *i.e.* Trajectron++ and our Joint- β -cVAE from

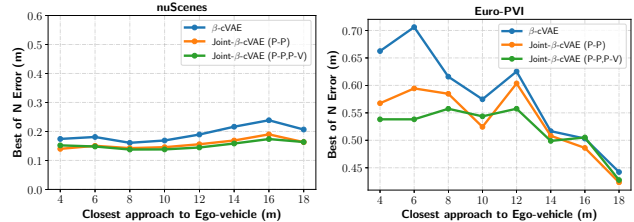


Figure 7. Error sorted by closest approach (proximity) to ego-vehicle. Higher error in close proximity to the ego-vehicle suggest dense interactions.

nuScenes (with both P-P,P-V interactions) to Euro-PVI in Table 4. We observe a considerable drop in performance in the Best of N error in comparison to the performance of the models when they are *both* trained and evaluated on Euro-PVI (Table 3). This provides additional evidence that the distribution of trajectories and interaction patterns in Euro-PVI is significantly different compared to nuScenes. We provide additional results in the supplemental material.

6. Conclusion and Outlook

We presented Euro-PVI, a new dataset with dense scenarios of vehicle-pedestrian (bicyclist) interaction and their trajectories to advance the task of future trajectory prediction which is crucial to the development of self-driving vehicles. We investigated the effect of interactions in urban environments on current state-of-the-art methods for existing nuScenes dataset which show a notable performance gap on our Euro-PVI dataset. To address this challenge of modeling complex interactions, we propose a Joint- β -cVAE approach. We demonstrate state of the art results both on nuScenes and on Euro-PVI. The performance advantage of our Joint- β -cVAE on Euro-PVI highlights the effectiveness of our approach in dense urban scenarios. The key to our success is a shared latent space between the interacting agents – which encodes the effect of intersections – in comparison to prior work which employ conditionally independent latent spaces. We believe that the Euro-PVI dataset along with Joint- β -CVAE approach provide a new important dimension to the task of future trajectory prediction with dense ego-vehicle - pedestrian (bicyclist) interactions.

References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Fei-Fei Li, and Silvio Savarese. Social LSTM: human trajectory prediction in crowded spaces. In *CVPR*, 2016.
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- [3] Apratim Bhattacharyya, Michael Hanselmann, Mario Fritz, Bernt Schiele, and Christoph-Nikolas Straehle. Conditional flow variational autoencoders for structured sequence prediction. In *NeurIPS Workshop*, 2019.
- [4] Apratim Bhattacharyya, Bernt Schiele, and Mario Fritz. Accurate and diverse sampling of sequences based on a "best of many" sample objective. In *CVPR*, 2018.
- [5] Apratim Bhattacharyya, Christoph-Nikolas Straehle, Mario Fritz, and Bernt Schiele. Haar wavelet based block autoregressive flows for trajectories. In *GCPR*, 2020.
- [6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.
- [7] Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Traphic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions. In *CVPR*, 2019.
- [8] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *CVPR*, 2019.
- [9] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [10] Nachiket Deo and Mohan M. Trivedi. Convolutional social pooling for vehicle trajectory prediction. In *CVPR Workshops*, 2018.
- [11] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *ICLR*, 2017.
- [12] Panna Felsen, Patrick Lucey, and Sujoy Ganguly. Where will they go? predicting fine-grained adversarial multi-agent motion using conditional variational autoencoders. In *ECCV*, 2018.
- [13] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *CVPR*, 2012.
- [14] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S. Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, Tiffany Fernandez, Martin Jänicke, Sudesh Mirashi, Chiragkumar Savani, Martin Sturm, Oleksandr Vorobiov, Martin Oelker, Sebastian Garreis, and Peter Schuberth. A2D2: audi autonomous driving dataset. *CoRR*, 2020.
- [15] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: socially acceptable trajectories with generative adversarial networks. In *CVPR*, 2018.
- [16] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. In *Physical review E*, 1995.
- [17] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017.
- [18] John Houston, Guido Zuidhof, Luca Bergamini, Yawei Ye, Ashesh Jain, Sammy Omari, Vladimir Iglovikov, and Peter Ondruska. One thousand and one hours: Self-driving motion prediction dataset. *CoRR*, 2020.
- [19] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *ICCV*, 2019.
- [20] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [21] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian D. Reid, Hamid Rezaatofghi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In *NeurIPS*, 2019.
- [22] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B. Choy, Philip H. S. Torr, and Manmohan Chandraker. DESIRE: distant future prediction in dynamic scenes with interacting agents. In *CVPR*, 2017.
- [23] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer Graphics Forum*, 2007.
- [24] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G. Hauptmann, and Li Fei-Fei. Peeking into the future: Predicting future person activities and locations in videos. In *CVPR*, 2019.
- [25] James Lucas, George Tucker, Roger B. Grosse, and Mohammad Norouzi. Don't blame the elbow! A linear VAE perspective on posterior collapse. In *NeurIPS*, 2019.
- [26] Yuexin Ma, Xinge Zhu, Sibozhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In *AAAI*, 2019.
- [27] Srikanth Malla, Behzad Dariush, and Chiho Choi. TITAN: future forecast using action priors. In *CVPR*, 2020.
- [28] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien Gaidon. It is not the journey but the destination: Endpoint conditioned trajectory prediction. In *ECCV*, 2020.
- [29] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *CoRR*, 2014.
- [30] Abdullallah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian G. Claudel. Social-stgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *CVPR*, 2020.
- [31] Seong Hyeon Park, Gyubok Lee, Manoj Bhat, Jimin Seo, Minseok Kang, Jonathan Francis, Ashwin R. Jadhav, Paul Pu Liang, and Louis-Philippe Morency. Diverse and admissible trajectory forecasting through multimodal context understanding. In *ECCV*, 2020.
- [32] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You'll never walk alone: Modeling social behavior for multi-target tracking. In *ICCV*, 2009.
- [33] Amir Rasouli, Iuliia Kotseruba, Toni Kunic, and John K. Tsotsos. PIE: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction. In *ICCV*, 2019.

- [34] Nicholas Rhinehart, Kris M. Kitani, and Paul Vernaza. r2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In *ECCV*, 2018.
- [35] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatofghi, and Silvio Savarese. Sophie: An attentive GAN for predicting paths compliant to social and physical constraints. In *CVPR*, 2019.
- [36] Amir Sadeghian, Ferdinand Legros, Maxime Voisin, Ricky Vesel, Alexandre Alahi, and Silvio Savarese. Car-net: Clairvoyant attentive recurrent network. In *ECCV*, 2018.
- [37] Tim Salzman, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control. In *ECCV*, 2020.
- [38] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *NeurIPS*, 2015.
- [39] Haoran Song, Wenchao Ding, Yuxuan Chen, Shaojie Shen, Michael Yu Wang, and Qifeng Chen. Pip: Planning-informed trajectory prediction for autonomous driving. In *ECCV*, 2020.
- [40] Luca Anthony Thiede and Pratik Prabhanjan Brahma. Analyzing the variety loss in the context of probabilistic trajectory prediction. In *ICCV*, 2019.
- [41] Jakub M. Tomczak and Max Welling. VAE with a vampprior. In *AISTATS*, 2018.
- [42] Jun Xie, Martin Kiefel, Ming-Ting Sun, and Andreas Geiger. Semantic instance annotation of street scenes by 3d to 2d label transfer. In *CVPR*, 2016.
- [43] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *ECCV*, 2020.
- [44] Lidan Zhang, Qi She, and Ping Guo. Stochastic trajectory prediction with social graph network. *CoRR*, 2019.
- [45] Tianyang Zhao, Yifei Xu, Mathew Monfort, Wongun Choi, Chris L. Baker, Yibiao Zhao, Yizhou Wang, and Ying Nian Wu. Multi-agent tensor fusion for contextual trajectory prediction. In *CVPR*, 2019.