

# Hierarchical Latent Structure for Multi-Modal Vehicle Trajectory Forecasting

Dooseop Choi\*<sup>1</sup> and KyoungWook Min<sup>1</sup>

Artificial Intelligence Research Laboratory, ETRI  
{d1024.choi,kwmin92}@etri.re.kr

**Abstract.** Variational autoencoder (VAE) has widely been utilized for modeling data distributions because it is theoretically elegant, easy to train, and has nice manifold representations. However, when applied to image reconstruction and synthesis tasks, VAE shows the limitation that the generated sample tends to be blurry. We observe that a similar problem, in which the generated trajectory is located between adjacent lanes, often arises in VAE-based trajectory forecasting models. To mitigate this problem, we introduce a hierarchical latent structure into the VAE-based forecasting model. Based on the assumption that the trajectory distribution can be approximated as a mixture of simple distributions (or modes), the low-level latent variable is employed to model each mode of the mixture and the high-level latent variable is employed to represent the weights for the modes. To model each mode accurately, we condition the low-level latent variable using two lane-level context vectors computed in novel ways, one corresponds to vehicle-lane interaction and the other to vehicle-vehicle interaction. The context vectors are also used to model the weights via the proposed mode selection network. To evaluate our forecasting model, we use two large-scale real-world datasets. Experimental results show that our model is not only capable of generating clear multi-modal trajectory distributions but also outperforms the state-of-the-art (SOTA) models in terms of prediction accuracy. Our code is available at <https://github.com/d1024choi/HLSTrajForecast>.

## 1 Introduction

Trajectory forecasting has long been a great interest in autonomous driving since accurate predictions of future trajectories of traffic agents are essential for the safe motion planning of an autonomous vehicle (AV). Many approaches have been proposed for trajectory forecasting in the literature and remarkable progress has been made in recent years. The recent trend in trajectory forecasting is to predict multiple possible trajectories for each agent in the traffic scene. This is because human drivers' future behavior is uncertain, and consequently, the future motion of the agent naturally exhibits a multi-modal distribution.

Latent variable models, such as variational autoencoders (VAEs) [18] and generative adversarial networks (GANs) [13], have been used for modeling the

---

\* corresponding author

distribution over the agents’ future trajectories. Using latent variables, trajectory forecasting models can learn to capture agent-agent and agent-space interactions from data, and consequently, generate future trajectories that are compliant with the input scene contexts.

VAEs have been applied in many machine learning applications, including image synthesis [15, 32], language modeling [3, 33], and trajectory forecasting [5, 20] because they are theoretically elegant, easy to train, and have nice manifold representations. One of the limitations of VAEs is that the generated sample tends to be blurry (especially in image reconstruction and synthesis tasks) [36]. We found from our experiments that a similar problem often arises in VAE-based trajectory forecasting models. More specifically, it is often found that the generated trajectory is located between adjacent lanes as illustrated in Figure 1. These false positive motion forecasts can cause uncomfortable rides for the AV with plenty of sudden brakes and steering changes [6]. In the rest of this paper, we will refer to this problem as *mode blur* as instance-level lanes are closely related to the modes of the trajectory distribution [16]. Mode blur is also found in the recent SOTA model [8] as shown in supplementary materials.

Many approaches have been proposed to mitigate the blurry sample generation problem primarily for image reconstruction or synthesis tasks. In this paper, we introduce a hierarchical latent structure into a VAE-based forecasting model to mitigate mode blur. Based on the assumption that the trajectory distribution can be approximated as a mixture of simple distributions (or modes), the low-level latent variable is employed to model each mode of the mixture and the high-level latent variable is employed to represent the weights for the modes. As a result, the forecasting model is capable of generating clear multi-modal trajectory distributions. To model each mode accurately, we condition the low-level latent variable using two lane-level context vectors (one corresponds to vehicle-lane interaction (VLI) and the other to vehicle-vehicle interaction (V2I)) computed in novel ways. The context vectors are also used to model the weights via the proposed mode selection network. Lastly, we also introduce two techniques to further improve the prediction performance of our model: 1) positional data preprocessing and 2) GAN-based regularization. The preprocessing is introduced based on the fact that vehicles moving along a lane usually try to be parallel to the tangent vector of the lane. The regularization is intended to ensure that the proposed model generates trajectories that match the shape of the lanes well.

In summary, our contributions are the followings:

- The hierarchical latent structure is introduced in the VAE-based forecasting model to mitigate mode blur.
- Two context vectors (one corresponds to the VLI and the other to the V2I) calculated in novel ways are proposed for lane-level scene contexts.
- Positional data preprocessing and GAN-based regularization are introduced to further improve the prediction performance.
- Our forecasting model outperforms the SOTA models in terms of prediction accuracy on two large-scale real-world datasets.

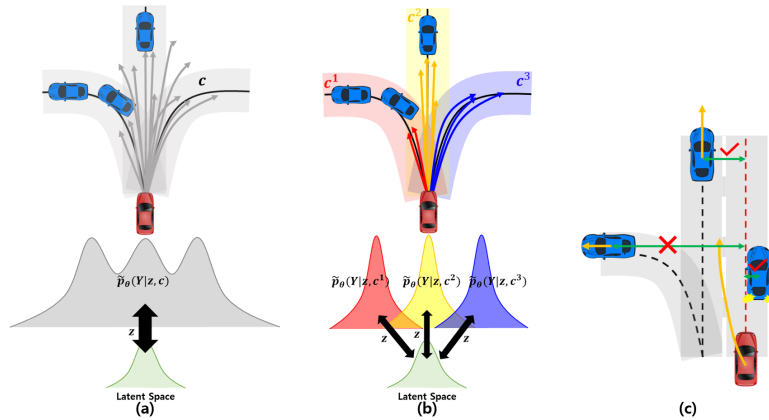


Fig. 1: Mode blur in trajectory forecasting and our approach. (a) Generated trajectories tend to locate between adjacent lanes. (b) We let a latent variable model each mode of the trajectory distribution to mitigate mode blur. (c) The target vehicle (red) takes into account not only its reference lane (red dashed line) but also the surrounding lanes (black dashed lines) and the surrounding vehicles. Only the surrounding vehicles within a certain distance from the reference lanes (green lines with arrows) influences the future motion of the target vehicle.

## 2 Related Works

### 2.1 Limitations of VAEs

The VAE framework has been used to explicitly learn data distributions. The models based on the VAE framework learn mappings from samples in a dataset to points in a latent space and generate plausible samples from variables drawn from the latent space. The VAE-based generative models are known to suffer from two problems: 1) posterior collapse (that the models ignore the latent variable when generating samples) and 2) blurry sample generation. To mitigate the problems, many approaches have been proposed in the literature, primarily for image reconstruction or synthesis tasks [11, 14, 15, 17, 27, 28, 32, 35]. In trajectory forecasting, some researchers [5, 30] have employed the techniques for the mitigation of the posterior collapse. To mitigate the blurry sample generation, [2] proposed a “best-of-many” sample objective that leads to accurate and diverse trajectory generation.

### 2.2 Forecasting with Lane Geometry

Because the movement of vehicles on the road is greatly restricted by the lane geometry, many works have been proposed to utilize the lane information provided by High-Definition (HD) maps [5, 9, 10, 12, 16, 22, 23, 25, 26, 30]. There are two types of approaches to the representation of the lane information: 1) rasterizing the components of the HD maps on a 2D canvas to obtain the top-view

images of the HD maps, 2) representing each component of the HD maps as a series of coordinates of points. In general, Convolutional Neural Network (CNN) is utilized for the former case while Long Short-Term Memory (LSTM) or 1D-CNN is utilized for the latter case to encode the lane information. In this paper, we adopt the second approach. The centerline of each lane in the HD maps is first represented as a series of equally-spaced 2D coordinates and then encoded by an LSTM network. The ability to handle individual lanes in the HD maps allows us to calculate lane-level scene contexts.

### 2.3 Lane-level Scene Context

Since instance-level lanes are closely related to the modes of the trajectory distribution, recent works [10, 16, 23, 25] proposed calculating lane-level scene contexts and using them for generating trajectories. Our work shares the idea with the previous works. However, ours differs from them in the way it calculates the lane-level scene contexts, which leads to significant gains in the prediction performance. Instead of considering only a single lane for a lane-level scene context, we also take into account surrounding lanes along with their relative importance. The relative importance is calculated based on the past motion of the target vehicle, thus reflecting the vehicle-lane interaction. In addition, for the interaction between the target vehicle and surrounding vehicles, we consider only the surrounding vehicles within a certain distance from the reference lane as illustrated in Figure 1c. This approach shows improved prediction performance compared to the existing approaches that consider either all neighbors [25] or only the most relevant neighbor [16]. This result is consistent with the observation that only a subset of surrounding vehicles is indeed relevant when predicting the future trajectory of the target vehicle [21].

## 3 Proposed Method

In this section, we present the details of our trajectory forecasting model.

### 3.1 Problem Formulation

Assume that there are  $N$  vehicles in the traffic scene. We aim to generate plausible trajectory distributions  $p(\mathbf{Y}_i|\mathbf{X}_i, \mathcal{C}_i)$  for the vehicles  $\{V_i\}_{i=1}^N$ . Here,  $\mathbf{X}_i = \mathbf{p}_i^{(t-H:t)}$  denotes the positional history of  $V_i$  for the previous  $H$  timesteps at time  $t$ ,  $\mathbf{Y}_i = \mathbf{p}_i^{(t+1:t+T)}$  denotes the future positions of  $V_i$  for the next  $T$  timesteps, and  $\mathcal{C}_i$  denotes additional scene information available to  $V_i$ . For  $\mathcal{C}_i$ , we use the positional histories of the surrounding vehicles  $\{\mathbf{X}_j\}_{j=1, j \neq i}^N$  and the lane candidates  $\mathbf{L}^{(1:M)}$  available for  $V_i$  at time  $t$ , where  $\mathbf{L}^m = \mathbf{I}_{1, \dots, F}^m$  denotes the  $F$  equally spaced coordinate points on the centerline of the  $m$ -th lane. Finally, we note that every positional information is expressed in the coordinate frame

defined by  $V_i$ 's current position and heading. According to [16],  $p(\mathbf{Y}_i|\mathbf{X}_i, \mathcal{C}_i)$  can be re-written as

$$p(\mathbf{Y}_i|\mathbf{X}_i, \mathcal{C}_i) = \sum_{m=1}^M \underbrace{p(\mathbf{Y}_i|E_m, \mathbf{X}_i, \mathcal{C}_i)}_{\text{mode}} \underbrace{p(E_m|\mathbf{X}_i, \mathcal{C}_i)}_{\text{weight}}, \quad (1)$$

where  $E_m$  denotes the event that  $\mathbf{L}^m$  becomes the reference lane for  $V_i$ . Equation 1 shows that the trajectory distribution can be expressed as a weighted sum of the distributions which we call *modes*. The fact that the modes are usually much simpler than the overall distribution inspired us to model each mode through a latent variable, and sample trajectories from the modes in proportion to their weights as illustrated in Figure 1b.

### 3.2 Forecasting Model with Hierarchical Latent Structure

We introduce two latent variables  $\mathbf{z}_l \in \mathbb{R}^D$  and  $\mathbf{z}_h \in \mathbb{R}^M$  to model the modes and the weights for the modes in Eq. 1. With the low-level latent variable  $\mathbf{z}_l$ , our forecasting model defines  $p(\mathbf{Y}_i|E_m, \mathbf{X}_i, \mathcal{C}_i)$  by using the decoder network  $p_\theta(\mathbf{Y}_i|\mathbf{z}_l, \mathbf{X}_i, \mathcal{C}_i^m)$  and the prior network  $p_\gamma(\mathbf{z}_l|\mathbf{X}_i, \mathcal{C}_i^m)$  based on

$$p(\mathbf{Y}_i|E_m, \mathbf{X}_i, \mathcal{C}_i) = \int_{\mathbf{z}_l} p(\mathbf{Y}_i|\mathbf{z}_l, \mathbf{X}_i, \mathcal{C}_i^m) p(\mathbf{z}_l|\mathbf{X}_i, \mathcal{C}_i^m) d\mathbf{z}_l, \quad (2)$$

where  $\mathcal{C}_i^m \subset \mathcal{C}_i$  denotes the scene information relevant to  $\mathbf{L}^m$ . To train our forecasting model, we employ the conditional VAE framework [31] and optimize the following modified ELBO objective [14]:

$$\begin{aligned} \mathcal{L}_{ELBO} = & -\mathbb{E}_{\mathbf{z}_l \sim q_\phi} [\log p_\theta(\mathbf{Y}_i|\mathbf{z}_l, \mathbf{X}_i, \mathcal{C}_i^m)] \\ & + \beta KL(q_\phi(\mathbf{z}_l|\mathbf{Y}_i, \mathbf{X}_i, \mathcal{C}_i^m) || p_\gamma(\mathbf{z}_l|\mathbf{X}_i, \mathcal{C}_i^m)), \end{aligned} \quad (3)$$

where  $\beta$  is a constant and  $q_\phi(\mathbf{z}_l|\mathbf{Y}_i, \mathbf{X}_i, \mathcal{C}_i^m)$  is the approximated posterior network. The weights for the modes  $p(E_m|\mathbf{X}_i, \mathcal{C}_i)$  are modeled by the high-level latent variable  $\mathbf{z}_h$ , which is output of the proposed mode selection network  $\mathbf{z}_h = f_\varphi(\mathbf{X}_i, \mathcal{C}_i^{(1:M)})$ .

As shown in Eq. 3 and the definition of the mode selection network, the performance of our forecasting model is dependent on how the lane-level scene information  $\mathcal{C}_i^m$  is utilized along with  $\mathbf{X}_i$  for defining the lane-level scene context. One can consider two interactions for the lane-level scene context: the VLI and V2I. This is because the future motion of the vehicle is highly restricted not only by the vehicle's motion history but also by the motion histories of the surrounding vehicles and the lane geometry of the road. For the VLI, the existing works [10, 16, 23, 25] considered only the reference lane. For the V2I, [16] considered only one vehicle most relevant to the reference lane, while the others considered all vehicles. In this paper, we present novel ways of defining the two interactions. For the VLI, instead of considering only the reference lane, we also take into account surrounding lanes along with their relative importance, which

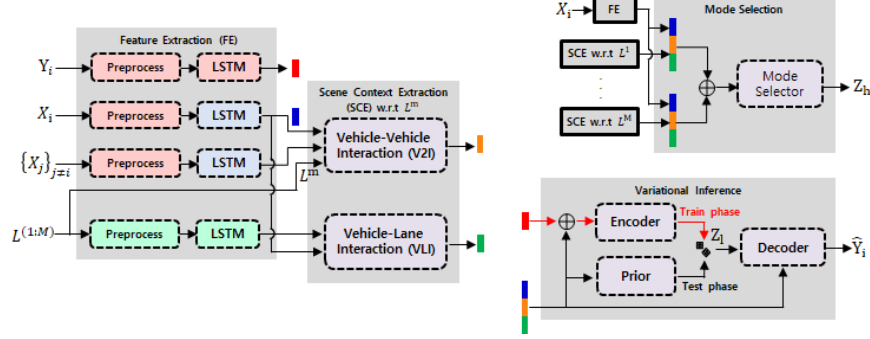


Fig. 2: Overall architecture of our forecasting model. To generate  $K$  future trajectories of  $V_i$ , lane-level scene context vectors  $\{c_i^m\}_{m=1}^M$ , each of which corresponds to one of  $\mathbf{L}^{(1:M)}$ , are first calculated via scene context extraction module. Next,  $\{w_m\}_{m=1}^M$  ( $w_m$  denotes the probability that  $V_i$  will drive along  $\mathbf{L}^m$  in the future) are calculated by using  $z_h$ . Finally,  $\lfloor K \times w_m \rfloor$  out of  $K$  future trajectories are generated by the decoder network using  $c_i^m$  and  $z_l$ .

is calculated based on the target vehicle’s motion history. The V2I is encoded through a GNN by considering only surrounding vehicles within a certain distance from the reference lane. Our approach is based on the fact that human drivers often pay attention to surrounding lanes and vehicles occupying the surrounding lanes when driving along the reference lane. Driving behaviors such as lane changes and overtaking are examples.

### 3.3 Proposed Network Structure

We show in Fig. 2 the overall architecture of our forecasting model. In the following sections, we describe the details of our model.

**Feature Extraction Module:** Three LSTM networks are used to encode the positional data  $\{\mathbf{X}_a\}_{a=1}^N$ ,  $\mathbf{Y}_i$ , and  $\mathbf{L}^{(1:M)}$ , respectively. The last hidden state vector of the networks is used for the encoding result. Before the encoding process, we preprocess the positional data. For the vehicles, we calculate the speed and heading at each timestep and concatenate the sequential speed and heading data to the original data along the data dimension. As a result,  $\{\mathbf{X}_a\}_{a=1}^N$  and  $\mathbf{Y}_i$  have the data dimension of size 4 (x-position, y-position, speed, and heading). For the lanes, at each coordinate point, we calculate the tangent vector and the direction of the tangent vector. The sequential tangential and directional data are concatenated to the original data along the data dimension. As a result,  $\mathbf{L}^{(1:M)}$  have the data dimension of size 5 (2D position vector, 2D tangent vector, and direction). We introduce the preprocessing step to make our model better infer the future positions of the target vehicle with the historical speed and heading records and the tangential data, based on that vehicles moving along a lane

usually try to be parallel to the tangent vector of the lane. As shown in Table 1, the prediction performance of our model is improved due to the preprocessing step. In the rest of this paper, we use a tilde symbol at the top of a variable to indicate that it is the result of the encoding process. For example, the encoding result of  $\mathbf{X}_i$  is expressed as  $\tilde{\mathbf{X}}_i$ .

**Scene Context Extraction Module:** Two lane-level context vectors are calculated in this stage. Assume that  $L^m$  is the reference lane for  $V_i$ . The context vector  $\mathbf{a}_i^m$  for the VLI is calculated as follows:

$$\mathbf{a}_i^m = [\tilde{\mathbf{L}}^m; \sum_{l=1, l \neq m}^M \alpha_l \tilde{\mathbf{L}}^l], \quad (4)$$

where  $\{\alpha_l\}_{l=1}^M$  are the weights calculated through the attention operation [1] between  $\tilde{\mathbf{X}}_i$  and  $\tilde{\mathbf{L}}^{(1:M)}$  and the semi-colon denotes the concatenation operation.  $\alpha_l$  represents the relative importance of the surrounding lane  $L^l$  compared to the reference lane under the consideration of the past motion of  $V_i$ . As a result, our model can generate plausible trajectories for the vehicles that drive paying attention to multiple lanes. For example, suppose that the vehicle is changing its lane from  $L^m$  to  $L^l$ .  $\alpha_l$  will be close to 1 and  $\mathbf{a}_i^m$  can be approximated as  $[\tilde{\mathbf{L}}^m; \tilde{\mathbf{L}}^l]$ , thus, our model can generate plausible trajectories corresponding to the lane change. We show in supplementary materials how the target vehicle interacts with the surrounding lanes of the reference lane using some driving scenarios.

To model the interaction between  $V_i$  and its surrounding vehicles  $\{V_j\}_{j \neq i}$ , we use a GNN. As we mentioned, only the surrounding vehicles within a certain distance from the reference lane are considered for the interaction; see Fig. 1c. Let  $\mathcal{N}_i^m$  denote the set of the vehicles including  $V_i$  and its select neighbors. The context vector  $\mathbf{b}_i^m$  for the V2I is calculated as follows:

$$\mathbf{m}_{j \rightarrow i} = \text{MLP}([\mathbf{p}_j^t - \mathbf{p}_i^t; \mathbf{h}_i^k; \mathbf{h}_j^k]), \quad (5)$$

$$\mathbf{o}_i = \sum_{j \in \mathcal{N}_i^m, j \neq i} \mathbf{m}_{j \rightarrow i}, \quad (6)$$

$$\mathbf{h}_i^{k+1} = \text{GRU}(\mathbf{o}_i, \mathbf{h}_i^k), \quad (7)$$

$$\mathbf{b}_i^m = \sum_{j \in \mathcal{N}_i^m, j \neq i} \mathbf{h}_j^{K-1}, \quad (8)$$

where  $\mathbf{h}^0 = \tilde{\mathbf{X}}$  for all vehicles in  $\mathcal{N}_i^m$ . The message passing from  $V_j$  to  $V_i$  is defined in Eq. 5 and all messages coming to  $V_i$  are aggregated by the sum operation as shown in Eq. 6. After the  $K$  rounds of the message passing, the hidden feature vector  $\mathbf{h}_j^{K-1}$  represents not only the motion history of  $V_j$  but also the history of the interaction between  $V_j$  and the others. The distance threshold  $\tau$  for  $\mathcal{N}_i^m$  plays the important role in the performance improvement. We explore the choice

of  $\tau$  value and empirically find that the best performance is achieved with  $\tau = 5$  meters (the distance between two nearby lane centerlines in straight roads is around 5 meters). Finally, note that we use the zero vector for  $\mathbf{b}_i^m$  when  $\mathcal{N}_i^m$  has the target vehicle only.

**Mode Selection Network:** The weights for the modes of the trajectory distribution are calculated by the mode selection network  $\mathbf{z}_h = f_\varphi(\mathbf{X}_i, \mathcal{C}_i^{(1:M)})$ . As instance-level lanes are closely related to the modes, it can be assumed that there are  $M$  modes, each corresponding to one of  $\mathbf{L}^{(1:M)}$ . We calculate the weights from the lane-level scene context vectors  $\mathbf{c}_i^m = [\tilde{\mathbf{X}}_i; \mathbf{a}_i^m; \mathbf{b}_i^m]$  which condense the information about the modes:

$$\mathbf{z}_h = \text{MLP}_{f_\varphi}([\mathbf{c}_i^1; \dots; \mathbf{c}_i^M]) \in \mathbb{R}^M. \quad (9)$$

The softmax operation is applied to  $\mathbf{z}_h$  to get the final weights  $\{w_m\}_{m=1}^M$ . Let  $\mathbf{z}_h^{SM}$  denote the result of applying the softmax operation to  $\mathbf{z}_h$ .  $w_m$  is equal to the  $m$ -th element of  $\mathbf{z}_h^{SM}$ . The lane-level scene context vector is the core feature vector for our encoder, prior, and decoder networks as described in the next section.

**Encoder, Prior, and Decoder:** The approximated posterior  $q_\phi(\mathbf{z}_l | \mathbf{Y}_i, \mathbf{X}_i, \mathcal{C}_i^m)$ , also known as encoder or recognition network, is implemented as MLPs with the encoding of the future trajectory and the lane-level scene context vector as inputs:

$$\mu_e, \sigma_e = \text{MLP}_{q_\phi}([\tilde{\mathbf{Y}}_i; \mathbf{c}_i^m]), \quad (10)$$

where  $\mu_e$  and  $\sigma_e$  are the mean and standard deviation vectors, respectively. The encoder is utilized in the training phase only because  $\mathbf{Y}_i$  is not available in the inference phase. The prior  $p_\gamma(\mathbf{z}_l | \mathbf{X}_i, \mathcal{C}_i^m)$  is also implemented as MLPs with the context vector as input:

$$\mu_p, \sigma_p = \text{MLP}_{p_\gamma}(\mathbf{c}_i^m), \quad (11)$$

where  $\mu_p$  and  $\sigma_p$  are the mean and standard deviation vectors, respectively. The latent variable  $\mathbf{z}_l$  is sampled from  $(\mu_e, \sigma_e)$  via the re-parameterization trick [18] during the training and from  $(\mu_p, \sigma_p)$  during the inference.

The decoder network generates the prediction of the future trajectory,  $\hat{\mathbf{Y}}_i$ , via an LSTM network as follows:

$$\mathbf{e}_i^t = \text{MLP}_{emb}(\hat{\mathbf{p}}_i^t), \quad (12)$$

$$\mathbf{h}_i^{t+1} = \text{LSTM}([\mathbf{e}_i^t; \mathbf{c}_i^m; \mathbf{z}_l], \mathbf{h}_i^t), \quad (13)$$

$$\hat{\mathbf{p}}_i^{t+1} = \text{MLP}_{dec}(\mathbf{h}_i^{t+1}), \quad (14)$$

where we initialize  $\hat{\mathbf{p}}_i^0$  and  $\mathbf{h}_i^0$  as the last observed position of  $V_i$  and the zero-vector, respectively.



### 3.4 Regularization Through GAN

To generate more clear image samples, [19] proposed a method that combines VAE and GAN. Based on the observation that the discriminator network implicitly learns a rich similarity metric for images, the typical element-wise reconstruction metric (e.g.,  $L_2$ -distance) in the ELBO objective is replaced with a feature-wise metric expressed in the discriminator. In this paper, we also propose training our forecasting model with a discriminator network simultaneously. However, we don't replace the element-wise reconstruction metric with the feature-wise metric since the characteristic of trajectory data is quite different from that of images. We instead use the discriminator to regularize our forecasting model during the training so that the trajectories generated by our model well match the shape of the reference lane.

The proposed discriminator network is defined as follows:

$$s = D(\mathbf{Y}_i, \mathbf{L}^m) = \text{MLP}_{dis}([\tilde{\mathbf{Y}}_i; \tilde{\mathbf{L}}^m]) \in \mathbb{R}^1. \quad (15)$$

We explored different choices for the encoding of the inputs to the discriminator network and observed that the following approaches improve the prediction performance: 1)  $\tilde{\mathbf{Y}}_i$  is the result of encoding  $[\mathbf{Y}_i; \Delta\mathbf{Y}_i]$  through an LSTM network where  $\Delta\mathbf{Y}_i = \Delta\mathbf{p}_i^{(t+1:t+T)}$ ,  $\Delta\mathbf{p}_i^t = \mathbf{p}_i^t - \mathbf{l}_f^m$ , and  $\mathbf{l}_f^m$  is the coordinate point of  $\mathbf{L}^m$  closest to  $\mathbf{p}_i^t$ , 2)  $\tilde{\mathbf{L}}^m$  is from the feature extraction module. We also observed that generating trajectories for the GAN objective ( $\mathcal{L}_{GAN}$  defined in Eq. 18) from both the encoder and prior yields better prediction performance, which is consistent with the observations in [19]. However, not back-propagating the error signal from the GAN objective to the encoder and prior does not lead to the performance improvement, which is not consistent with the observations in [19].

### 3.5 Training Details

The proposed model is trained by optimizing the following objective:

$$\mathcal{L} = \mathcal{L}_{ELBO} + \alpha\mathcal{L}_{BCE} + \kappa\mathcal{L}_{GAN}. \quad (16)$$

Here,  $\mathcal{L}_{BCE}$  is the binary cross entropy loss for the mode selection network and is defined as follows:

$$\mathcal{L}_{BCE} = \text{BCE}(\mathbf{g}^m, \text{softmax}(\mathbf{z}_h)), \quad (17)$$

where  $\mathbf{g}^m$  is the one-hot vector indicating the index of the lane, in which the target vehicle traveled in the future timesteps, among the  $M$  candidate lanes.  $\mathcal{L}_{GAN}$  is the typical adversarial loss defined as follows:

$$\mathcal{L}_{GAN} = \mathbb{E}_{\mathbf{Y} \sim p_{data}}[\log D(\mathbf{Y}, \mathbf{L})] + \mathbb{E}_{\mathbf{z} \sim p_z}[\log(1 - D(G(\mathbf{z}), \mathbf{L}))], \quad (18)$$

where  $G$  denotes our forecasting model. The hyper-parameters ( $\alpha$ ,  $\kappa$ ) in Eq. 16 and  $\beta$  in Eq. 3 are set to 1, 0.01, and 0.5, respectively. More details can be found in supplementary materials.

### 3.6 Inference

Future trajectories for the target vehicle are generated from the modes based on their weights. Assume that  $K$  trajectories need to be generated for  $V_i$ .  $\lfloor K \times w_m \rfloor$  out of  $K$  future trajectories are generated by the decoder network using  $\mathbf{c}_i^m$  and  $\mathbf{z}_l$ . In the end, a total of  $K$  trajectories can be generated from  $\{\mathbf{c}_i^m\}_{m=1}^M$  since  $\sum_{m=1}^M w_m = 1$ .

## 4 Experiments

### 4.1 Dataset

Two large-scale real-world datasets, Argoverse Forecasting [7] and nuScenes [4], are used to evaluate the prediction performance of our model. Both provide 2D or 3D annotations of road agents, track IDs of agents, and HD map data. nuScenes includes 1000 scenes, each 20 seconds in length. A 6-second future trajectory is predicted from a 2-second past trajectory for each target vehicle. Argoverse Forecasting is the dataset for the trajectory prediction task. It provides more than 300K scenarios, each 5 seconds in length. A 3-second future trajectory is predicted from a 2-second past trajectory for each target vehicle. Argoverse Forecasting and nuScenes publicly release only training and validation sets. Following the existing works [16, 30], we use the validation set for the test. For the training, we use the training set only.

### 4.2 Evaluation Metric

For the quantitative evaluation of our forecasting model, we employ two popular metrics, average displacement error (ADE) and final displacement error (FDE), defined as follows:

$$ADE(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{p}}^t - \mathbf{p}^t\|_2, \quad (19)$$

$$FDE(\hat{\mathbf{Y}}, \mathbf{Y}) = \|\hat{\mathbf{p}}^T - \mathbf{p}^T\|_2, \quad (20)$$

where  $\mathbf{Y}$  and  $\hat{\mathbf{Y}}$  respectively denote the ground-truth trajectory and its prediction. In the rest of this paper, we denote  $ADE_K$  and  $FDE_K$  as the minimum of ADE and FDE among the  $K$  generated trajectories, respectively. It is worth noting that  $ADE_1$  and  $FDE_1$  metrics shown in the tables presented in the later sections represent the average quality of the trajectories generated for  $\mathbf{Y}$ . Our derivation can be found in the supplementary materials. On the other hand,  $ADE_K$  and  $FDE_K$  represent the quality of the trajectory closest to the ground-truth among the  $K$  generated trajectories. We will call  $ADE_{K \geq 12}$  and  $FDE_{K \geq 12}$  metrics in the tables the *best quality* in the rest of this paper. According to [5], the average quality and the best quality are complementary and evaluate the precision and coverage of the predicted trajectory distributions, respectively.

Model	PDP	VLI	V2I	GAN	$ADE_1/FDE_1$	$ADE_{15}/FDE_{15}$
<b>M1</b>	✗	✗	✗	✗	3.15/7.53	0.95/1.82
<b>M2</b>	✓	✗	✗	✗	3.03/7.22	0.93/1.80
<b>M3</b>	✓	✓	✗	✗	2.91/7.00	0.94/1.82
<b>M4</b>	✓	✓	✓	✗	2.67/6.38	0.91/1.77
<b>M5</b>	✓	✓	✓	✓	<b>2.64/6.32</b>	<b>0.89/1.72</b>

(a)

Model	$ADE_1/FDE_1$	$ADE_{15}/FDE_{15}$
<b>Ours (<math>\tau=1</math>)</b>	2.66/ <b>6.31</b>	0.92/1.76
<b>Ours (<math>\tau=5</math>)</b>	<b>2.64/6.32</b>	<b>0.89/1.72</b>
<b>Ours (<math>\tau=10</math>)</b>	2.65/6.34	0.95/1.84
<b>Ours+All</b>	2.67/6.34	1.00/1.98
<b>Ours+Rel</b>	2.75/6.52	0.91/1.77

(b)

Table 1: Ablation study conducted on nuScenes




Fig. 3: Mode blur example

### 4.3 Ablation Study

**Performance Gain over Baseline** In Table 1a, we present the contributions of each idea to the performance gain over a baseline. **M1** denotes the baseline that does not use the positional data preprocessing (PDP), VLI, V2I, and GAN regularization proposed in this paper. We can see from the table that the average quality of the generated trajectories is improved by both the PDP and the VLI (**M1** v.s. **M2** v.s. **M3**). The improvement due to the VLI is consistent with the observation in [16] that consideration of multiple lane candidates is more helpful than using a single best lane candidate in predicting the future trajectory. Both the average quality and the best quality are much improved by the V2I (**M3** v.s. **M4**). The accurate trajectory prediction for the vehicles waiting for traffic lights is the most representative case of the performance improvement by the V2I. Due to the past movement of the neighboring vehicles waiting for the traffic light, our model can easily conclude that the target vehicle will also be waiting for the traffic light. Finally, the prediction performance is further improved by the GAN regularization (**M4** v.s. **M5**). As seen in Eq. 15, our discriminator uses a future trajectory along with the reference lane to discriminate between fake trajectories and real trajectories.

**Effect of Surrounding Vehicle Selection Mechanism** In Table 1b, we show the effect of the surrounding vehicle selection mechanism on the prediction

Model	$ADE_1$	$FDE_1$	$ADE_5$	$FDE_5$	$ADE_{10}$	$FDE_{10}$	$ADE_{15}$	$FDE_{15}$
CoverNet [26]	3.87	9.26	1.96	-	1.48	-	-	-
Trajectron++ [30]	-	9.52	1.88	-	1.51	-	-	-
AgentFormer [34]	-	-	1.86	3.89	1.45	2.86	-	-
ALAN [25]	4.67	10.0	1.77	<u>3.32</u>	<u>1.10</u>	<b>1.66</b>	-	-
LaPred [16]	<u>3.51</u>	<u>8.12</u>	<u>1.53</u>	3.37	1.12	2.39	1.10	2.34
MHA-JAM [24]	3.69	8.57	1.81	3.72	1.24	2.21	<u>1.03</u>	<b>1.7</b>
Ours	<b>2.64</b> <sub>0.87↓</sub>	<b>6.32</b> <sub>1.8↓</sub>	<b>1.33</b> <sub>0.2↓</sub>	<b>2.92</b> <sub>0.4↓</sub>	<b>1.04</b> <sub>0.06↓</sub>	<u>2.15</u> <sub>0.49↑</sub>	<b>0.89</b> <sub>0.14↓</sub>	<u>1.72</u> <sub>0.02↑</sub>

Table 2: Quantitative comparison on nuScenes

Model	$ADE_1$	$FDE_1$	$ADE_5$	$FDE_5$	$ADE_6$	$FDE_6$	$ADE_{12}$	$FDE_{12}$
DESIRE [20]	2.38	4.64	1.17	2.06	1.09	1.89	0.90	1.45
R2P2 [29]	3.02	5.41	1.49	2.54	1.40	2.35	1.11	1.77
VectorNet [12]	1.66	3.67	-	-	-	-	-	-
LaneAttention [23]	<u>1.46</u>	<u>3.27</u>	-	-	1.05	2.06	-	-
LaPred [16]	1.48	3.29	<u>0.76</u>	<u>1.55</u>	<u>0.71</u>	<u>1.44</u>	<u>0.60</u>	<u>1.15</u>
Ours	<b>1.44</b> <sub>0.02↓</sub>	<b>3.15</b> <sub>0.12↓</sub>	<b>0.70</b> <sub>0.06↓</sub>	<b>1.35</b> <sub>0.2↓</sub>	<b>0.65</b> <sub>0.06↓</sub>	<b>1.24</b> <sub>0.2↓</sub>	<b>0.51</b> <sub>0.09↓</sub>	<b>0.85</b> <sub>0.3↓</sub>

Table 3: Quantitative comparison on Argoverse Forecasting

performance of our model. Here, **Ours** ( $\tau$ ) denotes our model in which only the surrounding vehicles within  $\tau$  meters from the reference lane are considered. **Ours+Rel** and **Ours+All** denote our model in which the most relevant vehicle and all the vehicles are considered, respectively. We can see from the table that **Ours** with  $\tau = 5$  shows the best performance. This result demonstrates that considering only surrounding vehicles within a certain distance from the reference lane is effective in modeling the V2I from a lane-level perspective.

**Hierarchical Latent Structure** We show in Fig. 3 the generated trajectories for a particular scenario to demonstrate how helpful the introduction of the hierarchical latent structure would be for the mitigation of mode blur. In the figure, **Baseline** denotes the VAE-based forecasting model in which a latent variable is trained to model the trajectory distribution. **Baseline+BOM** and **Baseline+NF** respectively denote **Baseline** trained with the best-of-many (BOM) sample objective [2] and normalizing flows (NF) [28]. We introduce NF since the blurry sample generation is often attributed to the limited capability of the approximated posterior [15] and NF is a powerful framework for building flexible approximated posterior distributions [17]. In the figure, gray and black circles indicate historical and future positions, respectively. Squares with colors indicate the predictions of the future positions. Time is encoded in the rainbow color map ranging from red (0s) to blue (6s). Red solid lines indicate the centerlines of the candidate lanes. For the scenario, fifteen trajectories were generated. We can see in the figure that the proposed model generates trajectories that are aligned with the lane candidates. In contrast, neither normalizing flows nor BOM objective can help a lot for the mitigation of mode blur.

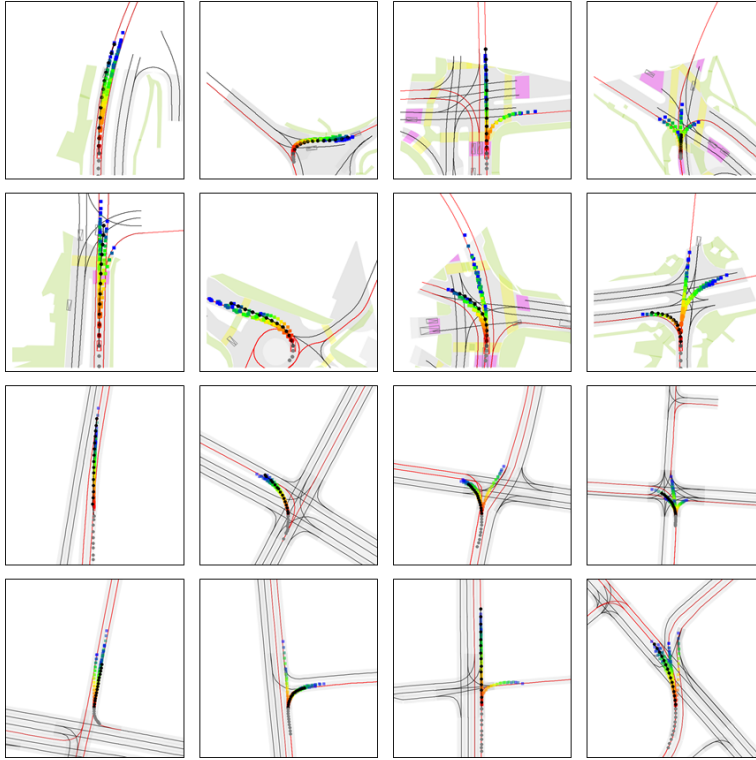


Fig. 4: Trajectory prediction examples of our forecasting model on nuScenes (the first and second rows) and Argoverse Forecasting (the third and fourth rows)

#### 4.4 Performance Evaluation

**Quantitative Evaluation** We compare our forecasting model with the existing models objectively. The results are shown in Table 2 and 3. Note that the bold and underline indicate the best and second-best performance, respectively. The values in the subscript indicate the performance gain over the second-best or loss over the best. Finally, the values in the table are from the corresponding papers and [16]. Table 2 presents the results on Nuscenes. It shows that our model outperforms the SOTA models [16, 25, 34] on most of the metrics. In particular, the performance gains over the SOTA models in the  $ADE_{K \leq 5}$  and  $FDE_{K \leq 5}$  metrics are significant. Consequently, it can be said that the trajectories generated from our model, on average, are more accurate than those from the SOTA models. On the other hand, [25] shows the significant performance on  $FDE_{10}$ . This is because, in [25], the vehicle trajectory is defined along the centerlines in a 2D curvilinear normal-tangential coordinate frame, so that the predicted trajectory is well aligned with the centerlines. However, [25] shows the poorest performance in the average quality. Table 3 presents the results on Argoverse Forecasting. It is seen that our forecasting model outperforms the SOTA models [16, 23] on

all the metrics. The  $ADE_{12}$  and  $FDE_{12}$  results show that our model achieves much better performance in the best quality compared to the models. However, the performance gain over the second-best model in the average quality is not significant. In short, our forecasting model exhibits remarkable performance in the average and best quality on the two large-scale real-world datasets.

**Qualitative Evaluation** Figure 4 illustrates the trajectories generated by our model for particular scenarios in the test dataset. Note that fifteen and twelve trajectories were generated for each scenario in nuScenes and Argoverse Forecasting, respectively. We can see in the figure that the generated trajectories are well distributed along admissible routes. In addition, the shape of the generated trajectory matches the shape of the candidate lane well. These results verify that the trajectory distribution is nicely modeled by the two latent variables conditioned by the proposed lane-level scene context vectors. It is noticeable that our model can generate plausible trajectories for the driving behaviors that require simultaneous consideration of multiple lanes. The first and third figures in the first column show the scenario where the target vehicle has just started changing lanes, and the second shows the scenario where the target vehicle is in the middle of a lane change. For both scenarios, our model generates plausible trajectories corresponding to both changing lanes and returning back to its lane. Finally, the last figure in the first column shows the scenario where the target vehicle is in the middle of a right turn. Our model well captures the motion ambiguity of the vehicle that can keep a lane or change lanes.

## 5 Conclusions

In this paper, we proposed a VAE-based trajectory forecasting model that exploits the hierarchical latent structure. The hierarchy in the latent space was introduced to the forecasting model to mitigate mode blur by modeling the modes of the trajectory distribution and the weights for the modes separately. For the accurate modeling of the modes and weights, we introduced two lane-level context vectors calculated in novel ways, one corresponds to the VLI and the other to the V2L. The prediction performance of the model was further improved by the two techniques, positional data preprocessing and GAN-based regularization, introduced in this paper. Our experiments on two large-scale real-world datasets demonstrated that the model is not only capable of generating clear multi-modal trajectory distributions but also outperforms the SOTA models in terms of prediction accuracy.

**Acknowledgment** This research work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIP) (No. 2020-0-00002, Development of standard SW platform-based autonomous driving technology to solve social problems of mobility and safety for public transport-marginalized communities)

## References

1. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. In: *Int. Conf. on Learn. Represent.* (2015)
2. Bhattacharyya, A., Schiele, B., Fritz, M.: Accurate and diverse sampling of sequences based on a best-of-many sample objective. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
3. Bowman, S.R., Vilnis, L., Vinyals, O., Dai, A.M., Jozefowicz, R., Bengio, S.: Generating sentences from a continuous space. In: *arXiv:1511.06349* (2015)
4. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: a multimodal dataset for autonomous driving. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2020)
5. Casas, S., Gulino, C., Suo, S., Luo, K., Liao, R., Urtasun, R.: Implicit latent variable model for scene-consistent motion forecasting. In: *Eur. Conf. Comput. Vis.* (2020)
6. Casas, S., Gulino, C., Suo, S., Urtasun, R.: The importance of prior knowledge in precise multimodal prediction. In: *Int. Conf. Intell. Robots Syst.* (2020)
7. Chang, M.F., Lambert, J., Sangkloy, P., Singh, J., Bak, S., Hartnett, A., Wang, D., Carr, P., Lucey, S., Ramanan, D., Hays, J.: Argoverse: 3d tracking and forecasting with rich maps. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2019)
8. Cui, A., Sadat, A., Casas, S., Liao, R., Urtasun, R.: Lookout: diverse multi-future prediction and planning for self-driving. In: *Int. Conf. Comput. Vis.* (2021)
9. Cui, H., Radosavljevic, V., F.-C.Chou, Lin, T.H., Nguyen, T., Huang, T.K., Schneider, J., Djuric, N.: Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In: *IEEE Int. Conf. Robotics and Automation* (2019)
10. Fang, L., Jiang, Q., Shi, J., Zhou, B.: TpNet: trajectory proposal network for motion prediction. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2020)
11. Fu, H., Li, C., Liu, X., Gao, J., Celikyilmaz, A., Carin, L.: Cyclical annealing schedule: A simple approach to mitigating kl vanishing. In: *NAACL* (2019)
12. Gao, J., Sun, C., Zhao, H., Shen, Y., Anguelov, D., Li, C., Schmid, C.: Vectornet: encoding hd maps and agent dynamics from vectorized representation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2020)
13. Goodfellow, I., Abadie, J.P., Mirza, M., Xu, B., Farley, D.W., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Adv. Neural Inform. Process. Syst.* (2014)
14. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-vae: learning basic visual concepts with a constrained variational framework. In: *Int. Conf. on Learn. Represent.* (2017)
15. Huang, H., Li, Z., He, R., Sun, Z., Tan, T.: Introvae: Introspective variational autoencoders for photographic image synthesis. In: *Adv. Neural Inform. Process. Syst.* (2018)
16. Kim, B., Park, S.H., Lee, S., Khoshimjonov, E., Kum, D., Kim, J., Kim, J.S., Choi, J.W.: Lapred: lane-aware prediction of multi-modal future trajectories of dynamic agents. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)
17. Kingma, D.P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., Welling, M.: Improved variational inference with inverse autoregressive flow. In: *Adv. Neural Inform. Process. Syst.* (2016)
18. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: *arXiv:1312.6114* (2013)

19. Larsen, A.B.L., Sonderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: *Int. Conf. on Learn. Represent.* (2016)
20. Lee, N., Choi, W., Vernaza, P., Choy, C.B., Torr, P.H.S., Chan, M.: Desire: Distant future prediction in dynamic scenes with interacting agents. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2017)
21. Li, J., Yang, F., Ma, H., Malla, S., Tomizuka, M., Choi, C.: Rain: reinforced hybrid attention inference network for motion forecasting. In: *Int. Conf. Comput. Vis.* (2021)
22. Liang, M., Yang, B., Hu, R., Chen, Y., Liao, R., Feng, S., Urtasun, R.: Learning lane graph representations for motion forecasting. In: *Eur. Conf. Comput. Vis.* (2020)
23. Luo, C., Sun, L., Dabiri, D., Yuille, A.: Probabilistic multi-modal trajectory prediction with lane attention for autonomous vehicles. In: *IEEE Conf. Intell. Robots Syst.* (2020)
24. Messaoud, K., Deo, N., Trivedi, M.M., Nashashibi, F.: Trajectory prediction for autonomous driving based on multi-head attention with joint agent-map representation. In: *arXiv:2005.02545* (2020)
25. Narayanan, S., Moslemi, R., Pittaluga, F., Liu, B., Chandraker, M.: Divide-and-conquer for lane-aware diverse trajectory prediction. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)
26. P-Minh, T., Grigore, E.C., Boulton, F.A., Beijbom, O., Wolff, E.M.: Covernet: multimodal behavior prediction using trajectory sets. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2020)
27. Razavi, A., Oord, A., Poole, B., Vinyals, O.: Preventing posterior collapse with delta-vaes. In: *Int. Conf. on Learn. Represent.* (2019)
28. Rezende, D.J., Mohamad, S.: Variational inference with normalizing flows. In: *Int. Conf. on Mach. Learn.* (2015)
29. Rhinehart, N., Kitani, K.M., Vernaza, P.: R2p2: a reparameterized pushforward policy for diverse, precise generative path forecasting. In: *Eur. Conf. Comput. Vis.* (2018)
30. Salzmann, T., Ivanovic, B., Chakravarty, P., Pavone, M.: Trajectron++: dynamically-feasible trajectory forecasting with heterogeneous data. In: *Eur. Conf. Comput. Vis.* (2020)
31. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: *Adv. Neural Inform. Process. Syst.* (2015)
32. Vahdat, A., Kautz, J.: Nvae: a deep hierarchical variational autoencoder. In: *Adv. Neural Inform. Process. Syst.* (2020)
33. Yang, Z., Hu, Z., Salakhutdinov, R., B.-Kirkpatrick, T.: Improved variational autoencoders for text modeling using dilated convolutions. In: *Int. Conf. on Mach. Learn.* (2017)
34. Yuan, Y., Weng, X., Ou, Y., Kitani, K.: Agentformer: agent-aware transformers for socio-temporal multi-agent forecasting. In: *arXiv:2103.14023* (2021)
35. Zhao, S., Song, J., Ermon, S.: Infovae: information maximizing variational autoencoders. In: *arXiv:1706.02262* (2017)
36. Zhao, S., Song, J., Ermon, S.: Towards a deeper understanding of variational autoencoding models. In: *arXiv:1702.08658v1* (2017)