

StretchBEV: Stretching Future Instance Prediction Spatially and Temporally

Adil Kaan Akan[✉] and Fatma Güney[✉]

KUIS AI Center, Koc University, Istanbul, Turkey

{kakan20, fguney}@ku.edu.tr

<https://kuis-ai.github.io/stretchbev>

Abstract. In self-driving, predicting future in terms of location and motion of all the agents around the vehicle is a crucial requirement for planning. Recently, a new joint formulation of perception and prediction has emerged by fusing rich sensory information perceived from multiple cameras into a compact bird’s-eye view representation to perform prediction. However, the quality of future predictions degrades over time while extending to longer time horizons due to multiple plausible predictions. In this work, we address this inherent uncertainty in future predictions with a stochastic temporal model. Our model learns temporal dynamics in a latent space through stochastic residual updates at each time step. By sampling from a learned distribution at each time step, we obtain more diverse future predictions that are also more accurate compared to previous work, especially stretching both spatially further regions in the scene and temporally over longer time horizons. Despite separate processing of each time step, our model is still efficient through decoupling of the learning of dynamics and the generation of future predictions.

1 Introduction

Future prediction with sequential visual data has been studied from different perspectives. Stochastic video prediction methods operate in the pixel space and learn to predict future frames conditioned on the previous frames. These methods achieve impressive results by modelling the uncertainty of the future with stochasticity, however, long-term predictions in real-world sequences tend to be quite blurry due to the complexity of directly predicting pixels. A more practical approach is to consider a compact representation that is tightly connected to the modalities required for the downstream application. In self-driving, understanding the 3D properties of the scene and the motion of other agents plays a key role in future predictions. The bird’s-eye view (BEV) representation meets these requirements by first fusing information from multiple cameras into a 3D point cloud and then projecting the points to the ground plane [30]. This leads to a compact representation where the future location and motion of multiple agents in the scene can be reliably predicted. In this paper, we explore the potential of stochastic future prediction for self-driving to produce admissible and diverse results in long sequences with an efficient and compact BEV representation.

Future prediction from the BEV representation has been recently proposed in FIERY [21]. The BEV representations of past frames are first related in a temporal model and then used to learn two distributions representing the present and the future. Based on sampling from one of these distributions depending on train or test time, various future modalities are predicted with a recurrent model. For planning, long term multiple future predictions are crucial, however, the predictions of FIERY degrade over longer time spans due to the limited representation capability of a single distribution for increasing diversity in longer predictions. Following the official implementation, the predictions do not differ significantly based on random samples but converge to the mean, therefore lack diversity. We start from the same BEV representation and predict the same output modalities to be comparable. Differently, instead of two distributions for the present and the future, we propose to learn time-dependent distributions by predicting a residual change at each time step to better capture long-term dynamics. Furthermore, we show that by sampling a random variable at each time step, we can increase the diversity of future predictions while still being accurate and efficient. For efficiency, we use a state-space model [29] instead of costly auto-regressive models.

The main idea behind the efficiency of the state-space model is to decouple the learning of dynamics and the generation of predictions [14]. We first learn a low dimensional latent space from the BEV representation to capture the dynamics and then learn to decode some output modalities from the predictions in that latent space. These output modalities show the location and the motion of the agents in the scene. We can train our dynamics model independently by learning to match the BEV representations of future frames that are predicted by our model to the ones that are extracted from a pre-trained BEV segmentation model [30]. Through residual updates to the latent space, our model can capture the changes to the BEV representation over time. This way, the only information we use from the future is the BEV representation of future frames. Another option is to encode the future modalities to predict and provide the model with this encoded representation to learn a future distribution [21]. We experiment with both options in this paper. While providing labels in the future distribution leads to better performance, learning the dynamics becomes dependent on the labels. From the BEV predictions, we train a decoder to predict the location and the motion of future instances in a supervised manner. These output modalities increase the interpretability of the predicted BEV representations that can be used for learning a driving policy in future work.

We present a stochastic future instance prediction method in BEV from multiple cameras. We formulate the prediction task as learning temporal dynamics through stochastic residual updates to a latent representation of the state at each time step. Our model can generate diverse predictions which are interpretable through supervised decoding of the predictions. Our proposed approach clearly outperforms the state of the art in various metrics evaluating the decoded predictions, especially by large margins in challenging cases of spatially far regions and temporally long spans. We also show increased diversity in the predictions.

2 Related Work

2.1 Stochastic Future Prediction

Stochastic future prediction has been mostly explored in the context of next frame prediction in videos. In stochastic video generation, the goal is to predict future frames conditioned on a few initial frames. Typically, the main focus is the diversity of future predictions with a large number of samples for future and the number of frames to be predicted is at least twice as many as the initial conditioning frames. Our work fits into the stochastic future prediction framework, producing long term, diverse predictions, however, we predict future in the BEV space instead of the noisy pixel space.

Most of the stochastic video prediction methods [12,3] use a recurrent neural network in an auto-regressive manner by feeding the generated predictions back to the model to predict future. The performance of auto-regressive methods can be improved by increasing the network capacity [35] or introducing a hierarchy into the latent variables [7], which also increase the complexity of these methods. Due to complexity of predicting pixels, another group of work moves away from the pixel space to the keypoints [28] or to the motion space by incorporating motion history [1]. Our proposed approach follows a similar strategy by performing future prediction in the BEV representation, but more efficiently by avoiding auto-regressive predictions.

Auto-regressive strategy requires encoding the predictions, leading to high computational cost and creates a tight coupling between the temporal dynamics and the generation process [16,31]. The state-space models (SSM) break this coupling by separating the learning of dynamics from the generation process, resulting in a computationally more efficient approach. Low-dimensional states still depend on previous states but not on predictions. Furthermore, learned states can be directly used in reinforcement learning [16] and can be interpreted [31], making SSMs particularly appealing for self-driving. Earlier SSMs with variational deep inference suffer from complicated inference schemes and typically target low-dimensional data [25,23]. An efficient training strategy with a temporal model based on residual updates is proposed for high-dimensional video prediction in the state of the art SSM [14]. We adapt a similar residual update strategy for predicting future BEV representations. We also experimentally show that the content variable for the static part of the scene is not as effective in the BEV space as it is in pixel space [14].

2.2 Future Prediction in Driving

The typical approach to the prediction problem in self-driving is to first perform detection and tracking, and then do the trajectory prediction [8,19]. In these methods, errors are propagated at each step. There are some recent methods [27,6,5,13] which directly address the prediction problem from the sensory input including LiDAR, radar, and other sensors. These methods also typically rely on an HD map of the environment. Due to high performance and efficiency

of end-to-end approach, we follow a similar approach for future prediction but using cameras only and without relying on HD maps.

Despite their efficiency, most of the previous work focus on the most likely prediction [6] or only models the uncertainty regarding the ego-vehicle’s trajectory [5,13]. The motion prediction methods which consider the behavior of all the agents in the scene typically assume a top-down rasterised representation as input, e.g. Argoverse setting [9]. Even then, multiple future prediction problem is typically addressed by generating a fixed number of predictions [15,26,2], for example by estimating the likelihood of multiple target locations [36,17]. There are some exceptions [34,32,22] which directly address the diversity aspect with a probabilistic framework. These works, especially the ones using a deep variational framework [34,32] are similar to our approach in spirit, however, they operate in the coordinate space by assuming the availability of a top-down map where locations of agents are marked. We aim to learn this top-down BEV representation from multiple cameras by also segmenting the agents in the scene.

FIERY [21] is the first to address probabilistic future prediction from multiple cameras. However, future predictions are limited both in terms of diversity and length considering the typical video prediction setting. We propose a probabilistic future prediction method that can generate diverse predictions extending to different temporal horizons with a stochastic temporal dynamics model.

3 Methodology

3.1 A Compact Representation for Future Prediction

Modern self-driving vehicles are typically equipped with multiple cameras observing the scene from multiple viewpoints. Placing cameras on the vehicle is cheap but processing information even from a single camera can be quite expensive. The traditional approach in computer vision is to extract low-level and semantic cues from these cameras and then fuse them into a holistic scene representation to perform prediction and planning. Recent success of end-to-end methods in driving has led to a rethinking of this approach. Furthermore, building and maintaining HD maps require a significant effort which is expensive and hard to scale. A better approach is to learn a geometrically consistent scene representation which can also mark the location, the motion, and even the semantics of the dynamic objects in the scene.

The bird’s-eye view (BEV) representation initially proposed in [30], takes image \mathbf{x}_t^i at time t from each camera $i \in \{1, \dots, 6\}$ and fuses them into a compact BEV representation \mathbf{s}_t . This is achieved by encoding each image and also by predicting a distribution over the possible depth values. The BEV features are obtained by weighting the encoded image features according to depth probabilities predicted. These features are first lifted to 3D by using known camera intrinsic and extrinsic parameters and then the height dimension is pooled over to project the features into the bird’s-eye view. This results in the state representation \mathbf{s}_t that we use for future prediction as explained next.

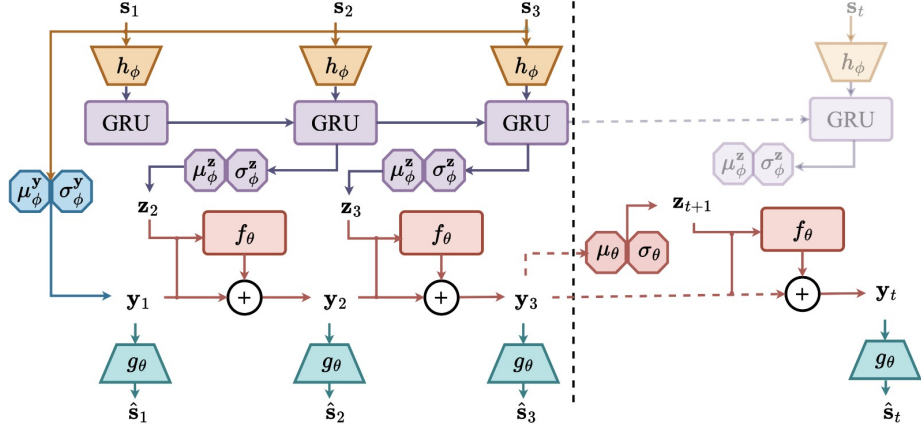


Fig. 1: **Architecture for Learning Temporal Dynamics.** This figure shows the inference procedure of our model StretchBEV. We start with the first $k = 3$ conditioning frames where we sample the stochastic latent variables from the posterior distribution (purple). On the right, we show the prediction at a step t after the conditioning frames where we sample from the learned future distribution (red). The dashed vertical line marks the conditioning frames.

3.2 Learning Temporal Dynamics in BEV

Notation In our formulation, $\mathbf{s}_{1:T}$ denotes a sequence of BEV feature maps representing the state of the vehicle and its environment for T time steps. In stochastic future prediction, the goal is to predict possible future states $\hat{\mathbf{s}}_{k+1:T}$ conditioned on the state in the first k time steps. Precisely, we condition on the first $k = 3$ steps and predict future in varying lengths from 4 to 12 steps ahead.

Differently from previous work on stochastic video prediction [12, 14], the BEV state \mathbf{s}_t that is input to our stochastic prediction framework is the intermediate representations in a high dimensional space rather than a video frame in the pixel space as explained in Section 3.1. Similarly, the predicted output $\hat{\mathbf{s}}_t$ represents the predictions of future in the same high dimensional space. We decode these high dimensional future predictions into various output modalities $\hat{\mathbf{o}}_t$ such as future instance segmentation and motion as explained in Section 3.4. While these modalities need to be trained in a supervised manner in contrast to typical self-supervised stochastic video prediction frameworks [12, 14], they provide interpretability which is critical in self-driving. Furthermore, using these modalities in the posterior in addition to the future state representations improves the results significantly as we show in our experiments.

Stochastic Residual Dynamics Following [14], we learn the changes in the states through time with stochastic residual updates to a sequence of latent variables. For each state \mathbf{s}_t , there is a corresponding latent variable \mathbf{y}_t generating it, independent of the previous states (Fig. 1). Each \mathbf{y}_{t+1} only depends on the

previous \mathbf{y}_t and a stochastic variable \mathbf{z}_{t+1} . The randomness is introduced by the stochastic latent variable \mathbf{z}_{t+1} which is sampled from a normal distribution learned from the previous state’s latent variable only:

$$\mathbf{z}_{t+1} \sim \mathcal{N}(\mu_\theta(\mathbf{y}_t), \sigma_\theta(\mathbf{y}_t) \mathbf{I}) \quad (1)$$

Given \mathbf{z}_{t+1} , the dependency between the latent variables \mathbf{y}_t and \mathbf{y}_{t+1} is deterministic through the residual update:

$$\mathbf{y}_{t+1} = \mathbf{y}_t + f_\theta(\mathbf{y}_t, \mathbf{z}_{t+1}) \quad (2)$$

where f_θ is a small CNN to learn the residual updates to \mathbf{y}_t . We learn the distribution of future states from the corresponding latent variable as a normal distribution with constant diagonal variance: $\hat{\mathbf{s}}_t \sim \mathcal{N}(g_\theta(\mathbf{y}_t))$. The first latent variable is inferred from the conditioning frames by assuming a standard Gaussian prior: $\mathbf{y}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

On the Content Variable: In video prediction, a common practice is to represent the static parts of the scene with a content variable which allows the model to focus on the moving parts. On the contrary to the state of the art in video prediction [14], the content variable does not improve the results in our case (see Supplementary), therefore we omit it in our formulation here. This can be attributed to the details in the background that are confusing for learning dynamics while operating in the pixel space, but in our case, most of these details are already suppressed in the BEV representation.

On Diversity: In contrast to a present and a future distribution in FIERY [20,21], there is a distribution learned at each time step in our model. This corresponds to sampling stochastic random variables at each time step as opposed to sampling once to represent all the future frames. This is the key property which allows our model to produce diverse predictions in long sequences. By sampling from a learned distribution at each time step, our model learns to represent the complex dynamics of future frames, even for predictions further away from the conditioning frames. Furthermore, our model does not need a separate temporal block for learning the dynamics prior to learning these distributions. The dynamics are learned through the temporal evolution of latent variables by also considering the randomness of the future predictions with stochastic random variables at each time step. This not only increases the diversity of predictions but also alleviates the need for a separate temporal block, e.g. with 3D convolutions. Note that our formulation is still efficient, almost the same inference time as FIERY (see Supplementary), because the latent variables are low dimensional and each state is generated independently.

Moreover, FIERY uses only a single vector of latent variables which is expanded to the spatial grid to generate futures states probabilistically. Therefore, it uses the same stochastic noise in all the coordinates of the grid. However, in our model, we have a separate random variable at each coordinate of the grid to model the uncertainty spatially as well. Training FIERY with a separate random variable at each location of the grid results in diverging loss values.

3.3 Variational Inference and Architecture

Following the generative process in [14], the joint probability of the BEV states $\mathbf{s}_{1:T}$, the output modalities \mathbf{o}_t , and the latent variables $\mathbf{z}_{1:T}$ and $\mathbf{y}_{1:T}$ is as follows:

$$p(\mathbf{s}_{1:T}, \mathbf{o}_{1:T}, \mathbf{z}_{2:T}, \mathbf{y}_{1:T}) = p(\mathbf{y}_1) \prod_{t=2}^T p(\mathbf{z}_t, \mathbf{y}_t | \mathbf{y}_{t-1}) \prod_{t=1}^T p(\mathbf{o}_t | \mathbf{s}_t) p(\mathbf{s}_t | \mathbf{y}_t) \quad (3)$$

$$p(\mathbf{z}_t, \mathbf{y}_t | \mathbf{y}_{t-1}) = p(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{z}_t) p(\mathbf{z}_t | \mathbf{y}_{t-1}) \quad (4)$$

The relationship between \mathbf{y}_t and \mathbf{y}_{t-1} in $p(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{z}_t)$ (4) is deterministic through the stochastic latent residual as formulated in (2). Similarly for the term $p(\mathbf{o}_t | \mathbf{s}_t)$ in (3), the output modalities \mathbf{o}_t is learned from \mathbf{s}_t with a deterministic decoder in a supervised manner (Section 3.4).

Our goal is to maximize the likelihood of the BEV states extracted from the frames (Section 3.1) and the corresponding output modalities $p(\mathbf{s}_{1:T}, \mathbf{o}_{1:T})$. For that purpose, we learn a deep variational inference model q parametrized by ϕ which is factorized as follows:

$$\begin{aligned} q_{Z,Y} &\triangleq q(\mathbf{z}_{2:T}, \mathbf{y}_{1:T} | \mathbf{s}_{1:T}, \mathbf{o}_{2:T}) \\ &= q(\mathbf{y}_1 | \mathbf{s}_{1:k}) \prod_{t=2}^T q(\mathbf{z}_t | \mathbf{s}_{1:t}, \mathbf{o}_{2:t}) q(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{z}_t) \end{aligned} \quad (5)$$

where $k = 3$ is the number of conditioning frames and $q(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{z}_t)$ is equal to $p(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{z}_t)$ with the residual update as explained above. We obtain two versions of our model by keeping (StretchBEV-P) or removing (StretchBEV) the dependency of \mathbf{z}_t on $\mathbf{o}_{1:t}$ in the posterior in (5). We refer the reader to Supplementary for the derivation of the following evidence lower bound (ELBO):

$$\begin{aligned} \log p(\mathbf{s}_{1:T}, \mathbf{o}_{1:T}) &\geq \mathcal{L}(\mathbf{s}_{1:T}, \mathbf{o}_{1:T}; \theta, \phi) \\ &\triangleq -D_{\text{KL}}(q(\mathbf{y}_1 | \mathbf{s}_{1:k}) || p(\mathbf{y}_1)) \\ &\quad + \mathbb{E}_{(\tilde{\mathbf{z}}_{2:T}, \tilde{\mathbf{y}}_{1:T}) \sim q_{Z,Y}} \left[\sum_{t=1}^T \log p(\mathbf{s}_t | \tilde{\mathbf{y}}_t) p(\mathbf{o}_t | \mathbf{s}_t) \right. \\ &\quad \left. - \sum_{t=2}^T D_{\text{KL}}(q(\mathbf{z}_t | \mathbf{s}_{1:t}, \mathbf{o}_{1:t}) || p(\mathbf{z}_t | \tilde{\mathbf{y}}_{t-1})) \right] \end{aligned} \quad (6)$$

where D_{KL} denotes the Kullback-Leibler (KL) divergence, θ and ϕ represent model and variational parameters, respectively. Following the common practice, we choose $q(\mathbf{y}_1 | \mathbf{s}_{1:k})$ and $q(\mathbf{z}_t | \mathbf{s}_{1:t}, \mathbf{o}_{1:t})$ to be factorized Gaussian for analytically computing the KL divergences and use the re-parametrization trick [24] to compute gradients through the inferred variables. Then, the resulting objective function is maximizing the ELBO as defined in (6) and minimizing the supervised losses for the output modalities (Section 3.4).

We provide a summary of the steps in our temporal model as shown in Fig. 1. We start by fusing the images \mathbf{x}_t^i at time t from each camera i into the BEV state \mathbf{s}_t at each time step as explained in Section 3.1.

1. The resulting BEV states are still in high resolution, $\mathbf{s}_t \in \mathbb{R}^{C \times H \times W}$ where $(H, W) = (200, 200)$. Therefore, we first process them with an encoder h_ϕ to reduce the spatial resolution to 50×50 .
2. The first latent variable \mathbf{y}_1 is inferred using a convolutional neural network on the first three encoded states.
3. The stochastic latent variable \mathbf{z}_t is inferred at each time step from the respective encoded state, using a recurrent neural network which is a combination of a ConvGRU and convolutional blocks.
4. The residual change in the dynamics is predicted with f_θ based on both the previous state dynamics \mathbf{y}_t and the stochastic latent variable \mathbf{z}_{t+1} and added to \mathbf{y}_t to obtain \mathbf{y}_{t+1} .
5. From each \mathbf{y}_t , the state $\hat{\mathbf{s}}_t$ is predicted in the original resolution with g_θ .
6. Finally, the output modalities $\hat{\mathbf{o}}_t$ are decoded from the state prediction $\hat{\mathbf{s}}_t$.

3.4 Decoding Future Predictions

Based on the predictions of the future states at each time step, we train a supervised decoder to output semantic segmentation, instance center, offsets, and future optical flow, similar to the previous work [11,21]. The decoding function is a deterministic neural network that can be trained either jointly with the dynamics or independently, e.g. later for interpretability. The output modalities show both the location and the motion of instances at each time step. The motion predicted as future flow is used to track instances. We use the same supervised loss functions for each modality as the FIERY [21]. Although the decoding of future predictions is not necessary for planning and control, for example when training a driving agent to act on predictions, these predictions provide interpretability and allow to compare the methods in terms of various metrics evaluating each modality.

4 Experiments

4.1 Dataset and Evaluation Setting

We evaluate the performance of the proposed approach and compare to the state of the art method, FIERY [21], on the nuScenes dataset [4]. On the nuScenes, there are 6 cameras with overlapping views which provide the ego vehicle with a complete view of its surroundings. The nuScenes dataset consists of 1000 scenes with 20 seconds long at 2 frames per second.

We first follow the training and the evaluation setting proposed in FIERY [21] for comparison by using 1.0 second of past context to predict 2.0 seconds of future context. Given the sampling rate of 2 frames per second, this setting corresponds to predicting 4 future frames conditioned on 3 past frames. We call this setting *short* in terms of temporal length and define two more settings for longer temporal predictions. All the models are trained to predict 2.0 seconds into the future and only the evaluation is changed to predict longer time steps. In the *mid* and

long settings, we double and triple the number of future frames to predict, i.e. 8 and 12, respectively, that corresponds to 4.0 and 6.0 seconds into the future. These settings are closer to the stochastic video prediction setup [12,14,1] where there are typically many more frames to predict than the conditioning frames for measuring diversity and the performance of the models further away from the conditioning frames. Note that *short* and *long* refer to temporal length in our evaluations as opposed spatial coverage as defined in the previous work [21]. We also evaluate in terms of spatial coverage but call it *near* (30m×30m) and *far* (100m ×100m) for clarity.

4.2 Training Details

Our models follow the input and output setting proposed in the previous work [21]. We process 6 camera images at a resolution of 224×480 pixels for each frame and construct the BEV state of size $200 \times 200 \times 64$. We further process the states into a smaller spatial resolution (50×50) for efficiency before learning the temporal dynamics but increase it back to the initial resolution afterwards. Given the predicted states, we use the same decoder architecture as the FIERY to decode the object centers, the segmentation masks, the instance offsets, and the future optical flow at a resolution of 200×200 pixels. We provide the details of the architecture in Supplementary and we will release the code upon publication.

In our approach, learning temporal dynamics and decoding output modalities are separated from each other. Therefore, we can pre-train the temporal dynamics part without using the labels for the output modalities. In pre-training, our objective is to learn to match the future states that are extracted using a pre-trained BEV segmentation model [30], conditioned on the past states. This approach is more similar to self-supervised stochastic video prediction methods [12,14]. Furthermore, this way, learning of temporal dynamics can be improved by using camera sequences only as input which can be easily collected in large quantities. Then, we fine-tune the temporal dynamics with a smaller learning rate (see Supplementary for details) while learning to decode the output modalities in a supervised manner. The alternative is to jointly train the temporal dynamics and supervised decoding without pre-training. We present the results of our model StretchBEV with and without pre-training.

StretchBEV does not use the labels (\mathbf{o}_t) for learning the temporal dynamics, it only uses them in the supervised loss to decode the output modalities. In our full model StretchBEV-P, we encode output modalities following FIERY and use them in the posterior for learning the temporal dynamics. During training, we sample the stochastic latent variables from the posterior and learn to minimize the difference between the posterior and the future distribution. During inference, we sample from the posterior in the conditioning frames and sample from the learned future distribution in the following steps as shown in Fig. 1.

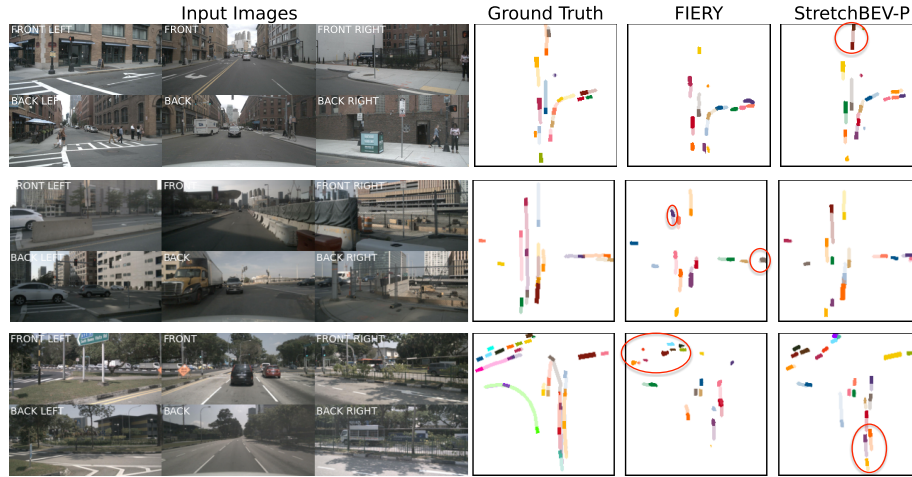


Fig. 2: **Qualitative Comparison over Different Temporal Horizons.** In this figure, we qualitatively compare the results of our model StretchBEV-P (**right**) to the ground truth (**left**) and FIERY [21] (**middle**) over short (**top**), mid (**middle**), and long (**bottom**) temporal horizons. Each color represents an instance of a vehicle with its trajectory trailing in the same color transparently.

4.3 Metrics

We use two different metrics for evaluating the decoded modalities, one frame level and another video level, that are also used in the previous work [21]. The first is Intersection over Union (IoU) to measure the quality of the segmentation at each frame. The second is Video Panoptic Quality (VPQ) to measure the quality of the segmentation and consistency of the instances through the video.

We evaluate the diversity quantitatively in terms of Generalized Energy Distance (D_{GED}) [33] by using $(1 - VPQ)$ as the distance as proposed in FIERY [21].

4.4 Ablation Study

In Table 1, we evaluate the effect of different versions of our model using IoU and VPQ metrics in the short temporal setting to be comparable to the previous work FIERY [21]. We reproduced their results as shown in the row *Reproduced*. In the first part of the table (StretchBEV), we show the results without explicitly using the labels for future prediction. In that case, labels are only used for decoding the output modalities and back-propagated to future prediction through decoding. Although this introduces a two-stage training, we believe that reporting results using this separation is important for future work to focus on future prediction with more unlabelled data. We measure the effect of pre-training by learning to match our future predictions to the results of a pre-trained model [30] in terms of the BEV state representation. Pre-training allows our model to learn the dynamics before decoding and improves the results significantly in each metric.

	Pre-training	Posterior w/labels	IoU (\uparrow)		VPQ (\uparrow)	
			Near	Far	Near	Far
StretchBEV	—	—	53.3	35.8	41.7	26.0
	✓	—	55.5	37.1	46.0	29.0
FIERY [21] Reproduced	—	✓	59.4	36.7	50.2	29.4
	—	✓	58.8	35.8	50.5	29.0
StretchBEV-P	—	✓	58.1	52.5	53.0	47.5

Table 1: **Ablation Study.** In this table, we present the results for the two versions of our model with (StretchBEV-P) and without (StretchBEV) using the labels for the output modalities in the posterior while learning the temporal dynamics and also show the effect of pre-training for the latter in comparison to FIERY [21] and our reproduced version of their results (Reproduced).

In the second half of the Table 1, we report the results using the labels in future prediction by explicitly feeding their encoding to the posterior distribution with the same encoding used in [21] to learn the future distribution. The difference between StretchBEV and StretchBEV-P is that the first has access to the BEV encoding of future predictions while the latter has access to both the BEV encoding and the encoding of the output modalities to predict in the posterior distribution. As can be seen from the results, both FIERY and our model using the labels in the future distribution perform better. This shows the importance of using a more direct and accurate information about future while learning the posterior. Compared to FIERY, our model can use the labels in the conditioning frames during inference and improves the results, especially in spatially far regions and in terms of VPQ, which point to a higher quality in our predictions stretching spatially and temporally over the video. Please see Supplementary for an extended version of Table 1 with multiple samples including standard deviation as an indication of uncertainty.

4.5 Temporally Long Predictions

In longer temporal horizons, future prediction becomes increasingly difficult. This is mainly due to increasing uncertainty of future further away from conditioning frames. In Fig. 3, we present the results over different temporal horizons for our model with pre-training without using the labels in the posterior (StretchBEV), FIERY [21], and our model by using the labels in the posterior (StretchBEV-P). There is a separate plot for IoU on the left and for VPQ on the right with respect to the future time steps predicted, ranging from 1s to 8s. The vertical line in 2s marks the training horizon. In Supplementary, we provide a table for the results over short, mid, and long temporal spans.

The negative effect of uncertain futures on each metric can be observed from the results of all the methods degrading from shorter to longer temporal spans. Our models perform better than FIERY in longer temporal spans. This is due




Fig. 3: **Evaluation over Different Temporal Horizons.** We plot the performance of our models StretchBEV and StretchBEV-P in comparison to FIERY [21] over a range of temporal horizons from 1 second to 8 seconds in terms of IoU (left) and VPQ (right) for spatially far (solid) and near (dashed) regions separately. The vertical dashed line marks the training horizon.

to better handling of uncertainty with stochastic latent residual variables. Our method StretchBEV-P outperforms FIERY by significant margins, especially in terms of far VPQ in longer temporal horizons, showing consistent predictions in the overall scene throughout the video. This can be attributed to the difficulty of locating small vehicles in spatially far regions. Since StretchBEV-P has access to the labels via the posterior in the conditioning frames, it learns the temporal dynamics to correctly propagate them to the future frames, while StretchBEV and FIERY struggle to locate the instances in the first place. FIERY learns a single distribution for present and future each, therefore we cannot utilize the labels in the conditioning frames with FIERY. The results of StretchBEV outperforming the other two methods in terms of near IOU in longer temporal spans is promising for future prediction methods with less supervision.

In Fig. 2, we qualitatively compare the performance of our model StretchBEV-P on the right to FIERY in the middle over short, mid, and long temporal horizons in each row. In the first row, our model predicts the future trajectories that are more similar to the ground truth shown on the left. For example, FIERY fails to predict the trajectory of the vehicle in front (marked with red circle). In the second row, our model correctly segments the vehicles, whereas FIERY misses several vehicles far on the right and also, predicts a vehicle that does not exist (in purple on the top left). In the third row, our model predicts the future trajectories of the vehicles correctly while FIERY misses some of the vehicles (marked with red circles). The challenging case of a vehicle turning on the left (green in ground truth) is missed by both models. Some of the vehicles are not visible on the input images, e.g. the back camera in the long temporal horizon. We provide a gif version of these results and more examples in Supplementary.

4.6 Segmentation

The previous work on bird’s-eye view segmentation typically focuses on single image segmentation task with a couple of exceptions focusing on prediction. In Table 2, we compare to two BEV segmentation prediction methods [18,21] using

Fishing-Cam [18]	Fishing-LiDAR [18]	FIERY [21]	StretchBEV	StretchBEV-P
30.0	44.3	57.3	<u>58.8</u>	65.7

Table 2: **Comparison of Semantic Segmentation Prediction.** In this table, we compare the predictions of our models, StretchBEV and StretchBEV-P for semantic segmentation to other BEV prediction methods in terms of IoU using the setting proposed in [18], i.e. $32.0\text{m} \times 19.2\text{m}$ at 10cm resolution over 2s future.

their setting with $32.0\text{m} \times 19.2\text{m}$ at 10cm resolution. Both methods predict 2s into the future which corresponds to our short temporal setting. FIERY [21] outperforms the previous method [18] even when using LiDAR, and our method significantly outperforms both methods.

4.7 Diversity

We quantitatively evaluate diversity by computing D_{GED} over 10 samples and show the results in Table 3 for our model StretchBEV-P and FIERY (our reproduced version). Our model outperforms FIERY with lower distance scores, demonstrating higher levels of diversity in the samples quantitatively. The difference is especially apparent in spatially far regions. For qualitative comparison, in Fig. 4, we visualize three samples from FIERY (left) and our model (right) over short, mid, and long temporal spans from top to bottom. While FIERY generates almost the same predictions in all three samples, our model can generate diverse predictions of future (marked with red). In the first row, our model can predict the turning behavior of the green vehicle at different speeds. In the second row, our model learns to adjust the speed of nearby vehicles proportionally, as in the case of purple, blue, and gray vehicles in the middle. Similarly, in the third row, our model can generate different predictions for the moving vehicles in the middle. Please see Supplementary for the gif results with more examples.

	Generalized Energy Distance (\downarrow)					
	Short		Mid		Long	
	Near	Far	Near	Far	Near	Far
FIERY [21]	106.09	140.36	118.74	147.26	127.18	152.38
StretchBEV	<u>103.97</u>	<u>132.38</u>	<u>114.11</u>	<u>138.15</u>	<u>119.01</u>	<u>142.51</u>
StretchBEV-P	82.04	85.51	94.02	98.45	101.90	109.12

Table 3: **Quantitative Evaluation of Diversity.** This table compares the results of our models to the reproduced results of FIERY [21] in terms of Generalized Energy Distance based on VPQ (lower better) for evaluating diversity.

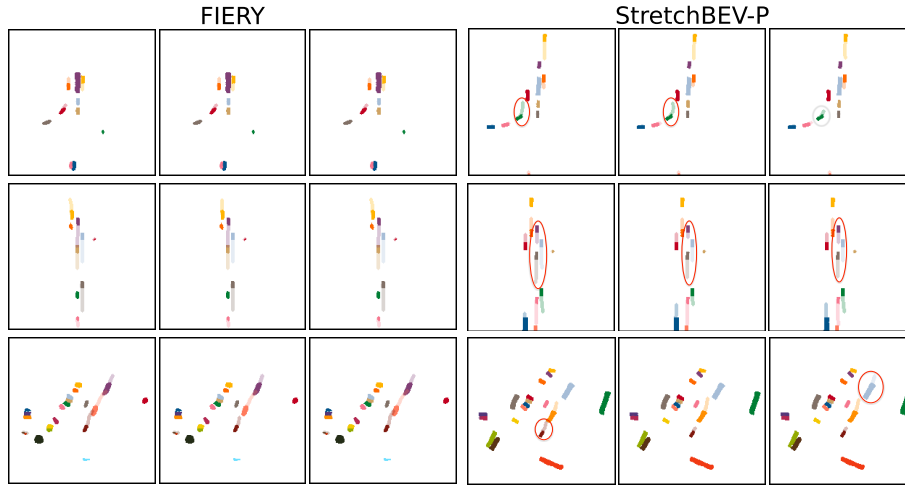


Fig. 4: **Qualitative Comparison of Diversity.** In this figure, we visualize random samples from FIERY [21] (left) and our model StretchBEV-P (right) over short (top), mid (middle), and long (bottom) temporal horizons.

5 Conclusion and Future Work

We introduced StretchBEV, a stochastic future instance prediction method that improve over the state of the art, especially in challenging cases, with more diverse predictions. We proposed two versions of our method with and without the labels for output modalities explicitly in the posterior while learning the dynamics. Both models improve the state of the art in spatially far regions and over temporally long horizons. Using labels in the posterior significantly improves the results in almost all metrics but introduces a dependency on the availability of labels in the conditioning frames during inference. Future work on learning dynamics should focus on closing the gap between the two approaches, for example with scheduled sampling.

Our temporal dynamics model can be interpreted as a Neural-ODE [10] because of its residual update dynamics. In our model, we use only one update in between time steps but in future, we plan to explore increasing the number of updates in between time step as done in the previous work [14]. We showed that our model increases the diversity of predictions due to improved modeling of stochasticity with sampling at every time step. In future, we plan to explore driving policies that can utilize stochastic future predictions. Learned latent states at each time step can be directly fed into a policy learning algorithm, e.g. as states in deep reinforcement learning. Furthermore, these states can be interpreted via supervised decoding into various future modalities that we predict.

Acknowledgments: K. Akan was supported by KUIS AI Center fellowship, F. Güney by TUBITAK 2232 International Fellowship for Outstanding Researchers.

References

1. Akan, A.K., Erdem, E., Erdem, A., Güney, F.: Slamp: Stochastic latent appearance and motion prediction. In: Proc. of the IEEE International Conf. on Computer Vision (ICCV) (2021)
2. Aydemir, G., Akan, A.K., Güney, F.: Trajectory forecasting on temporal graphs. arXiv preprint arXiv:2207.00255 (2022)
3. Babaeizadeh, M., Finn, C., Erhan, D., Campbell, R.H., Levine, S.: Stochastic variational video prediction. In: Proc. of the International Conf. on Learning Representations (ICLR) (2018)
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: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2020)
5. Casas, S., Gulino, C., Liao, R., Urtasun, R.: Spaggn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data. In: Proc. IEEE International Conf. on Robotics and Automation (ICRA) (2020)
6. Casas, S., Luo, W., Urtasun, R.: Intentnet: Learning to predict intention from raw sensor data. In: Proc. Conf. on Robot Learning (CoRL) (2018)
7. Castrejon, L., Ballas, N., Courville, A.: Improved conditional vrnn for video prediction. In: Proc. of the IEEE International Conf. on Computer Vision (ICCV) (2019)
8. Chai, Y., Sapp, B., Bansal, M., Anguelov, D.: Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In: Proc. Conf. on Robot Learning (CoRL) (2020)
9. Chang, M.F., Lambert, J.W., 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: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2019)
10. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. In: Advances in Neural Information Processing Systems (NeurIPS) (2018)
11. Cheng, B., Collins, M.D., Zhu, Y., Liu, T., Huang, T.S., Adam, H., Chen, L.C.: Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In: CVPR (2020)
12. Denton, E., Fergus, R.: Stochastic video generation with a learned prior. In: Proc. of the International Conf. on Machine learning (ICML) (2018)
13. Djuric, N., Cui, H., Su, Z., Wu, S., Wang, H., Chou, F., Martin, L.S., Feng, S., Hu, R., Xu, Y., Dayan, A., Zhang, S., Becker, B.C., Meyer, G.P., Vallespi-Gonzalez, C., Wellington, C.K.: Multixnet: Multiclass multistage multimodal motion prediction. In: Proc. IEEE Intelligent Vehicles Symposium (IV) (2021)
14. Franceschi, J.Y., Delasalles, E., Chen, M., Lamprier, S., Gallinari, P.: Stochastic latent residual video prediction. In: Proc. of the International Conf. on Machine learning (ICML) (2020)
15. 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: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2020)
16. Gregor, K., Besse, F.: Temporal difference variational auto-encoder. In: Proc. of the International Conf. on Learning Representations (ICLR) (2019)

17. Gu, J., Sun, C., Zhao, H.: Densetnt: End-to-end trajectory prediction from dense goal sets. In: Proc. of the IEEE International Conf. on Computer Vision (ICCV) (2021)
18. Hendy, N., Sloan, C., Tian, F., Duan, P., Charchut, N., Xie, Y., Wang, C., Philbin, J.: FISHING net: Future inference of semantic heatmaps in grids. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops (2020)
19. Hong, J., Sapp, B., Philbin, J.: Rules of the road: Predicting driving behavior with a convolutional model of semantic interactions. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2019)
20. Hu, A., Cotter, F., Mohan, N., Gurau, C., Kendall, A.: Probabilistic future prediction for video scene understanding. In: Proc. of the European Conf. on Computer Vision (ECCV) (2020)
21. Hu, A., Murez, Z., Mohan, N., Dudas, S., Hawke, J., Badrinarayanan, V., Cipolla, R., Kendall, A.: FIERY: Future instance segmentation in bird’s-eye view from surround monocular cameras. In: Proc. of the IEEE International Conf. on Computer Vision (ICCV) (2021)
22. Huang, X., McGill, S.G., DeCastro, J.A., Fletcher, L., Leonard, J.J., Williams, B.C., Rosman, G.: Diversitygan: Diversity-aware vehicle motion prediction via latent semantic sampling. *IEEE Robotics and Automation Letters (RA-L)* (2020)
23. Karl, M., Soelch, M., Bayer, J., van der Smagt, P.: Deep variational bayes filters: Unsupervised learning of state space models from raw data. In: Proc. of the International Conf. on Learning Representations (ICLR) (2017)
24. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: Proc. of the International Conf. on Learning Representations (ICLR) (2014)
25. Krishnan, R.G., Shalit, U., Sontag, D.: Structured inference networks for nonlinear state space models. In: Proc. of the Conf. on Artificial Intelligence (AAAI) (2017)
26. Liang, M., Yang, B., Hu, R., Chen, Y., Liao, R., Feng, S., Urtasun, R.: Learning lane graph representations for motion forecasting. In: Proc. of the European Conf. on Computer Vision (ECCV) (2020)
27. Luo, W., Yang, B., Urtasun, R.: Fast and furious: Real time end-to-end 3D detection, tracking and motion forecasting with a single convolutional net. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2018)
28. Minderer, M., Sun, C., Villegas, R., Cole, F., Murphy, K.P., Lee, H.: Unsupervised learning of object structure and dynamics from videos. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2019)
29. Murphy, K.P.: *Probabilistic Machine Learning: Advanced Topics*. MIT Press (2023), probml.ai
30. Pillion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: Proc. of the European Conf. on Computer Vision (ECCV) (2020)
31. Rubanova, Y., Chen, R.T.Q., Duvenaud, D.K.: Latent ordinary differential equations for irregularly-sampled time series. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2019)
32. Sriram, N.N., Liu, B., Pittaluga, F., Chandraker, M.: SMART: Simultaneous multi-agent recurrent trajectory prediction. In: Proc. of the European Conf. on Computer Vision (ECCV) (2020)
33. Székely, G.J., Rizzo, M.L.: The energy of data. *Annual Review of Statistics and Its Application* 4(1), 447–479 (2017)
34. Tang, Y.C., Salakhutdinov, R.: Multiple futures prediction. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2019)

35. Villegas, R., Pathak, A., Kannan, H., Erhan, D., Le, Q.V., Lee, H.: High fidelity video prediction with large stochastic recurrent neural networks. In: Advances in Neural Information Processing Systems (NeurIPS) (2019)
36. Zhao, H., Gao, J., Lan, T., Sun, C., Sapp, B., Varadarajan, B., Shen, Y., Shen, Y., Chai, Y., Schmid, C., Li, C., Anguelov, D.: TNT: target-driven trajectory prediction. In: Proc. Conf. on Robot Learning (CoRL) (2020)