A Framework for Probabilistic Generic Traffic Scene Prediction

Yeping Hu, Wei Zhan and Masayoshi Tomizuka

Abstract— In a given scenario, simultaneously and accurately predicting every possible interaction of traffic participants is an important capability for autonomous vehicles. The majority of current researches focused on the prediction of an single entity without incorporating the environment information. Although some approaches aimed to predict multiple vehicles, they either predicted each vehicle independently with no considerations on possible interaction with surrounding entities or generated discretized joint motions which cannot be directly used in decision making and motion planning for autonomous vehicle. In this paper, we present a probabilistic framework that is able to jointly predict continuous motions for multiple interacting road participants under any driving scenarios and is capable of forecasting the duration of each interaction, which can enhance the prediction performance and efficiency. The proposed traffic scene prediction framework contains two hierarchical modules: the upper module and the lower module. The upper module forecasts the intention of the predicted vehicle, while the lower module predicts motions for interacting scene entities. An exemplar real-world scenario is used to implement and examine the proposed framework.

I. INTRODUCTION

As the autonomous vehicle is becoming a big trend nowadays, safety is the most essential aspect to consider. Being able to predict future motions of the surrounding entities in the scene can greatly enhance the safety level of autonomous vehicles since potentially dangerous situations could be avoided in advance. Therefore, the Advanced Driver Assistance Systems (ADAS) is expected to have a full interpretation of the scene and be able to accurately predict possible behaviors of multiple traffic participants under various driving scenarios, which will then assure a safe, comfortable and cooperative driving experience. An illustration of the scene prediction is shown in Fig. 1.

There have been numerous works focusing on predicting the traffic participants in different driving scenarios. The predicted outcomes can be divided into two general categories : discrete and continuous. Intention estimation can be regarded as a discrete prediction problem, which is usually solved by classification strategies, such as Support Vector Machine (SVM) [1], Bayesian classifier [2], Hidden Markov Models (HMMs) [3], and Multilayer Perceptron (MLP) [4]. Motion prediction, on the other hand, is mostly treated as a regression problem, where it forecasts the short-term movements and long-term trajectories of vehicles. Various motion prediction methods use vehicle kinematic models at the prediction step and estimate the state recursively. For example, methods such as constant velocity (CV), constant acceleration (CA) and

Y. Hu, W. Zhan and M. Tomizuka are with the Department of Mechanical Engineering, University of California, Berkeley, CA 94720 USA [yeping_hu, wzhan, tomizuka@berkeley.edu]

Fig. 1. An illustration of the scene prediction results. The red car is the autonomous vehicle which is predicting the yellow car and its possible interaction with other scene entities.

Intelligent Driver Model (IDM) [5] have been wildly used. However, these methods are generally considered in simple traffic scenarios such as car following.

There are also many approaches that deal with motion predictions in more complicated situations such as lane changing and ramp merging. [6] utilized Artificial Neural Network (ANN) structure to predict vehicle lateral motions and used a SVM to further determine if a lane change will happen. A Variational Gaussian Mixture Model (VGMM) is proposed in [7] to classify and predict the long-term trajectories of vehicles in a simulated environment. The main limitation of these works, however, is that they failed to take surrounding vehicles into account, which is unreasonable since the trajectory of the predicted vehicle will be largely influenced by the environment.

To handle such problems, several works incorporated potentially affecting vehicles by making use of relational features. In [8], the authors used the long short-term memory (LSTM) to predict the most likely trajectories for vehicles in highway situation while considering their nine surrounding vehicles. [9] brought forward a Deep Neural Networks (DNN) to obtain the lateral acceleration and longitudinal velocity while taking into account five vehicles around the predicted car. Nevertheless, these methods only predict trajectories for a single scene entity without estimating the possible motions of other vehicles. Works such as [10] and [11] utilized a LSTM-based structure for each vehicle in the scene and predicted the probabilistic information on future locations over a occupancy grid map. Although possible trajectories for every surrounding vehicles were predicted, the authors treated each vehicle independently during the prediction process, which cannot provide sufficient and accurate predictions especially in highly interactive scenarios.

Very few studies have been done for simultaneously predicting multiple interacting traffic participants. The approaches of [12] and [13] do incorporate such interdependencies by jointly predicting behavior patterns of all on-

road vehicles. However, acquiring only the discrete behavior pattern is not enough for autonomous vehicles to fully predict the traffic scene or to directly perform risk assessment.

All the aforementioned motion prediction works did not have an estimation for the prediction horizon. They either fixed the prediction length beforehand or recursively estimated the vehicle state until a designated location is reached. However, for a given scene, multiple types of interaction between vehicles are possible and each of them is expected to have different time span. Therefore, it is not only irrational but also computationally expensive to predict every vehicle trajectories for the same time horizon.

In this paper, a probabilistic framework that is able to predict various types of dynamic scenes is proposed. It contains an upper module and a lower module, where the Semantic-based Intention and Motion Prediction (SIMP) method is used in the upper module and the Conditional Variational Autoencoder (CVAE) is used in the lower module. The upper module is capable of predicting possible semantic intention and motion of the selected vehicle, while the lower module can further predict joint probability distributions of motions for interacting traffic participants. Possible future trajectories will be sampled from each joint distribution and the prediction horizon for different interacting entities can be received from the upper module. Also, our framework can guarantee feasibilities for every sampled trajectory. These trajectories could then be easily used by the decision-making and motion planning process for autonomous vehicles.

The remainder of the paper is organized as follows: Section II provides the detailed explanation of the proposed scene prediction framework; Section III discusses an exemplar scenario to apply our framework; evaluations and results are provided in Section IV; and Section V concludes the paper.

II. SCENE PREDICTION FRAMEWORK

In this section, we first introduce the method applied to each of the two modules in the scene prediction framework. Then, the overall prediction process of the framework is illustrated.

A. Upper Module

For the upper module, we implemented the Semantic based Intention and Motion Prediction (SIMP) approach [14] due to its great adaptability for various scenarios and competitive prediction performance compared to other methods. It utilizes deep neural network to formulate a probabilistic framework which can predict the possible semantic intention and motion of the selected vehicle under various traffic scenes. The introduced semantics is defined as answering the question of *"Which area will the predicted vehicle most likely insert into? Where and when?"*. The inserted area is called Dynamic Insertion Area (DIA), which can be a available gap between any two vehicles on the road or can be a lane entrance/exit area. Each DIA is assigned to a 2D Gaussian Mixture Model (GMM) to predict a two dimensional vector: $\boldsymbol{y} = [y_s, y_t]^T$, where y_s describes vehicle's location and y_t

represents the time information. Note that both variables can be explicitly defined according to the problem formulation.

Therefore, given a set of input features $\boldsymbol{x}$, the probability distribution $\boldsymbol{y}_a$ over a single area a for the predicted vehicle can be expressed as:

$$p(\boldsymbol{y}_a|\boldsymbol{x}) = \sum_{m=1}^{M} \alpha_m \frac{1}{2\pi\sqrt{|\Sigma_m|}} exp\left(-\frac{D_m^T \Sigma_m^{-1} D_m}{2}\right), \quad (1)$$

where $D_m = \boldsymbol{y}_a - \mu_m$ and M denotes the total number of mixture components. For each mixture component m, the mixing coefficient α_m, mean μ_m, and covariance Σ_m formulate a probability density function of the output $\boldsymbol{y}_a$.

The output of the SIMP structure contains the required parameters for every 2D GMM and the weight w_a for each insertion area a. As for the desired outputs, not only the largest weight is expected to be associated to the actual inserted area, but also the highest probability is supposed to be at the proper location and time for the output distributions of that area. The loss function is then defined as

$$Loss = W_1\left(-\sum_{n} log\left\{\sum_{a=1}^{N_a} \hat{w}_a^n p(\boldsymbol{y}_a^n|\boldsymbol{x})\right\}\right)$$
$$+ W_2\left(-\sum_{n}\sum_{a=1}^{N_a} \hat{w}_a^n log(w_a^n)\right), \quad (2)$$

where N_a denotes the total number of DIA in the scene and $\hat{w}_a$ denotes the ground truth of the area weight, which is the one-hot-encoding of the final insertion area for the predicted vehicle. Parameters W1 and W2 are manually tuned such that the two loss components will have the same order of magnitude during training.

B. Lower Module

The lower module contains different motion models which will be assigned to each insertion area according to some criteria that will be discussed in Section III. Here, we applied the Conditional Variational Autoencoder (CVAE) [15][16] method to each motion model. CVAE has a similar structure to the typical variational autoencoder which contains an encoder and a decoder. It is rooted in bayesian inference, where the goal is to model the underlying probability distribution of the data so that new data could be sampled from that distribution. The overall structure of the CVAE we used is shown in Fig. 2.

In order to obtain the distribution of the output Y given the input X (i.e. $P(Y|X)$), a latent variable $z \sim \mathcal{N}(0, I)$ is introduced such that

$$P(Y|X) = \mathcal{N}(f(z, X), \sigma^2 \cdot I), \quad (3)$$

where it has mean f which is a function that can be directly learned from the data and covariance equal to the identity matrix I times a scalar σ.

Given the training data (X, Y), the framework first samples z from some arbitrary distribution Q different from

Fig. 2.   CVAE structure

$\mathcal{N}(0, I)$. According to Bayes rule, we have:

$$E_{z \sim Q}\big[\log P(Y|z, X)\big] = E_{z \sim Q}\big[\log P(z|Y, X) \\ + \log P(Y|X) - \log P(z|X)\big] \quad (4)$$

and by subtracting $E_{z \sim Q}(\log Q(z))$ from both sides:

$$\log P(Y|X) - E_{z \sim Q}\big[\log Q(z) - \log P(z|X, Y)\big] = \\ E_{z \sim Q}\big[\log P(Y|z, X) + \log P(z|X) - \log Q(z)\big]. \quad (5)$$

To make the right hand side closely approximates $\log P(Y|X)$, Q is constructed to depend on both X and Y, which will make $E_{z \sim Q}\big[\log Q(z) - \log P(z|X, Y)\big]$ small.

By writing the above expectation as Kullback-Leibler (KL) divergences, our variational objective is in the form:

$$\log P(Y|X) - D_{KL}\big[Q(z|X, Y)||P(z|X, Y)\big] = \\ E_{z \sim Q}\big[\log P(Y|z, X) - D_{KL}\big[Q(z|X, Y)||P(z|X)\big]\big], \quad (6)$$

where $P(z|X) = \mathcal{N}(0, I)$ since the model assumes that z is independent of X if Y is unknown. The right hand side can be optimized via stochastic gradient descent by using the reparameterization trick to enable the encoder to generate a vector of means and a vector of standard deviation. In general, Q is trained to "encode" Y into the latent z space such that the values of z can be "decoded" back to the output.

The loss function then becomes a summation of the generative loss, which is the mean square error between the network output and the ground truth, and a latent loss, which is the KL divergence term that forces the latent variables match a unit Gaussian:

$$Loss = ||Y - \hat{Y}||^2 + D_{KL}\big[Q(z|X, Y)||\mathcal{N}(0, I)\big], \quad (7)$$

where Y is the ground truth and $\hat{Y}$ is the output estimation. At test time, we can sample from the distribution $P(Y|X)$ by sampling z directly from $\mathcal{N}(0, I)$.

C. The Overall Framework

For a selected vehicle, its input feature vector x contains all the information of itself and its surrounding environment. The SIMP method takes x as the input and generates three types of outputs: w, $p(y_s|x)$, and $p(y_t|x)$. The vector w represents the weight for each of the N_a areas; $p(y_s|x)$ denotes the probabilistic distribution of the final destination within each DIA; and $p(y_t|x)$ denotes the probabilistic distribution of the time remained to enter each DIA for the predicted vehicle, which can be also interpreted as the time-to-lane-change (TTLC) distribution. These outputs along with the filtered input features will then feed into different motion models inside the lower module. Finally, the joint probability

---

Algorithm 1: Scene Prediction Framework

1 N_s : total number of output samples
2 f_s : data sampling rate
3 x : input vectors at current time
4 $w, p(y_s|x), p(y_t|x) \leftarrow SIMP(x)$
5 **for** $a = 1 : N_a$ **do**
6 | $M \leftarrow$ current motion model
7 | $n_a \leftarrow N_s \cdot w_a$
8 | **for** $i = 1 : n_a$ **do**
9 | | $T \leftarrow sample(p(y_{t,a}|x))$
10 | | $s_1, o_1 \leftarrow x$
11 | | **for** $j = 1 : (T \cdot f_s)$ **do**
12 | | | $T_j \leftarrow T - (j-1) \cdot f_s$
13 | | | $a_j \leftarrow sample(M(s_j, o_j, T_j))$
14 | | | **if** a_j *not feasible* **then**
15 | | | | redo iteration
16 | | | **else**
17 | | | | $s_{j+1} \leftarrow f(a_j, s_j)$
18 | | | **end**
19 | | **end**
20 | | **if** $p(y_{s,a} = s_{final}|x) < \epsilon$ **then**
21 | | | redo iteration
22 | | **end**
23 | **end**
24 **end**

---

distribution of motions over multiple interacting entities at the next time step can be obtained. At the current time step t, we denote states of the interacting vehicles as s_t, the other sensor observations of the surroundings as o_t, and the predicted actions as a_t. The next state of the interacting vehicles can be directly calculated using some mapping function f such that $s_{t+1} = f(a_t, s_t)$.

Therefore, given the current state information and according to the first-order Markov assumption, the joint probability distribution over the prediction horizon T can be expressed as:

$$p(s_{t+1}, s_{t+2}, ..., s_{t+T}|s_t) = p(s_{t+1}|s_t)p(s_{t+2}|s_{t+1}) \\ ...p(s_{t+T}|s_{t+T-1}), \quad (8)$$

where T has the distribution of $p(y_t|x)$. Since s_t is independent of o_t and T, we have:

$$p(s_{t+1}|s_t) = \sum_T \sum_{o_t} p(s_{t+1}, o_t, T|s_t) \\ = \sum_T \sum_{o_t} p(s_{t+1}|s_t, o_t, T)p(s_t), \quad (9)$$

where $p(s_{t+1}|s_t, o_t, T)$ is obtained from the output distribution of the motion prediction model and o_t is assumed to have a gaussian measurement noise. Here, we use sample-based method to infer the desired joint probability distribution and the overall process of the proposed scene prediction framework is illustrated in Algorithms 1.

Each motion model has its corresponding dynamic insertion area and the model will be triggered once its insertion
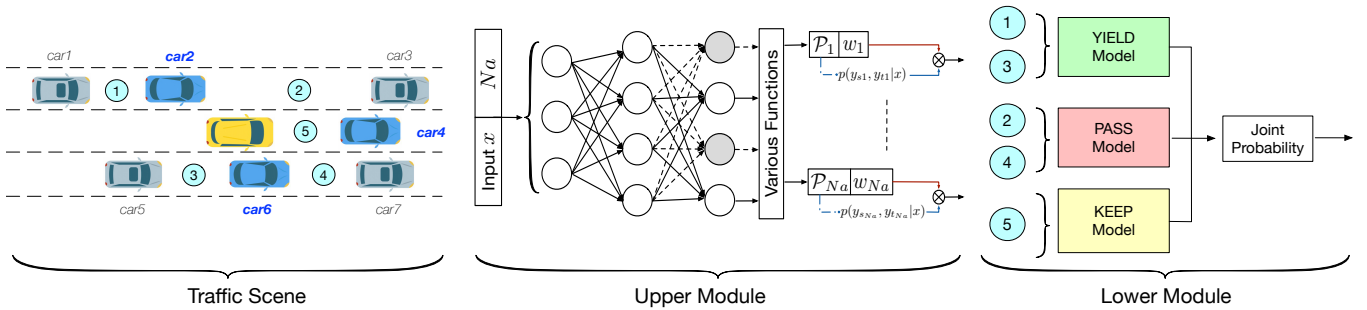
Fig. 3. The scene prediction framework for the selected exemplar scenario.

weight w_a is greater than a threshold (i.e. n_a is not zero). Within each trajectory sampling process, we first sample the time information T, from $p(y_t|\boldsymbol{x})$. Then, within the given prediction horizon $T \cdot f_s$, we iteratively apply our lower motion model to obtain the sampled output action, which will go through a feasibility check to make sure the predicted next state is reachable by the vehicle. Finally, we use the destination distribution $p(y_s|\boldsymbol{x})$ from the SIMP method to ensure our predicted final state is within the desired range.

III. AN EXEMPLAR SCENARIO

In this section, we use an exemplar highway driving scenario to apply our proposed scene prediction framework. The data source and detailed problem formulation are presented.

A. Dataset

The data we used was taken from the NGSIM dataset which is publicly available online at [17]. We used the US highway 101 dataset which contains detailed vehicle trajectory data collected with sampling rate $f_s = 10$ Hz. Since the original data, especially for the velocity and acceleration, are very noisy, we used an Extended Kalman Filter (EKF) for better estimations. For lane-change situation, we picked up to 40 frames (4s) before the vehicle's center intersects the lane mark; for lane-keep situation, 40 frames were selected for each vehicle. A total of 15,240 frames were chosen from the dataset and randomly split into 80% for training and 20% for testing.

B. Scenarios and Problem Description

The representation of the exemplar scenario and the corresponding framework structure are shown in Fig. 3. The yellow car is the predicted vehicle; the blue cars are the reference vehicles, which are vehicles that the predicted vehicle will most likely interact with; the grey cars are other surrounding vehicles that are assumed to be fully observable.

In the typical driving scenario, there are five circled areas (DIA) that our predicted vehicle could eventually enter. More specifically, if the predicted vehicle (yellow car) inserts into $area1$ or $area3$, it will yield to car2 or car6 respectively; if the predicted vehicle enters $area2$ or $area4$, it will pass car2 or car6 respectively; if, however, $area5$ is inserted, the predicted vehicle will keep its lane and follow car4. Therefore, we considered three motion models inside the

lower module: **yield**, **pass** and **keep**. Each model is trained separately on data that satisfy each particular model type. Note that for other traffic scenes such as ramp merging, the same type of motion models can be directly applied.

C. Features and Structure Details

The input features for the upper module are selected as same as in [14]. For the lower module, input states of the two interactive vehicles can be written as $\boldsymbol{s}_t = \{v_{pred}^t, v_{ref}^t, x_{ref}^t - x_{pred}^t, y_{ref}^t - y_{pred}^t\}$, where the subscript $pred$ denotes the predicted vehicle and ref denotes the reference vehicle, and the observed input $\boldsymbol{o}_t$ is determined by the type of the motion model. For example, $\boldsymbol{o}_t$ for the *pass* model contains states of the vehicle in front of the predicted and the reference car, whereas no $\boldsymbol{o}_t$ is needed for the *keep* model since only the front reference car need to be considered which information is already included in the input state. The predicted actions $\boldsymbol{a}_t$ for each motion model are expressed as $\boldsymbol{a}_t = \{\Delta x_{pred}^t, \Delta v_{pred}^t, \Delta x_{ref}^t, \Delta v_{ref}^t\}$, representing lateral displacements and the longitudinal velocity differences for both the predicted vehicle and the reference vehicle. Here, the aforementioned function f will map $\boldsymbol{a}_t$ and $\boldsymbol{s}_t$ linearly to $\boldsymbol{s}_{t+1}$.

For the SIMP structure, we use three fully connected layers of 400 neurons each, with $tanh$ non-linear activation functions. A dropout layer is appended to the end to enhance the network's generalization ability and prevent overfitting. For the CAVE structure, we utilize two fully connect layers of 128 neurons each for both the encoder and the decoder, and use three latent variables. The two modules are trained hierarchically and concatenated during testing.

IV. EVALUATION AND RESULTS

In this section, different evaluation methods are presented to assess the model quality and the final results are discussed.

A. Performance Evaluation

The output of the scene prediction framework are trajectories sampled from joint motion distributions of the predicted vehicle and each of its interacting entities (i.e. reference vehicles). We used the Root-Mean-Square Error (RMSE) as the validation metric to evaluate the output trajectories. However, multiple possible trajectories can be sampled from different motion models at the early stage of a lane changing
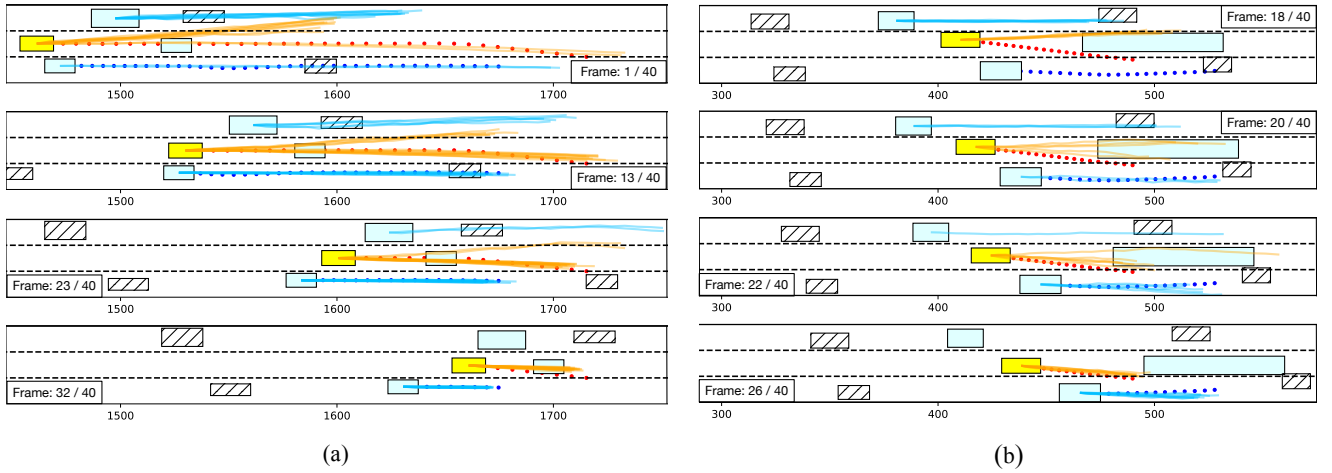
Fig. 4. Test results of two example cases. The yellow car represents the predicted vehicle; the blue cars are reference vehicles that might interact with the predicted vehicle; and the striped cars denote other surrounding vehicles which will not have direct interactions with the predict vehicle. The red and blue dotted lines are the ground truth trajectories (with 0.1s step-time) for the predicted vehicle and its interacting reference vehicle respectively. The solid lines are sampled trajectories from the motion models.

scenario and more trajectories will be sampled from the ground-truth motion model as the time-to-lane-change value gets smaller. Therefore, in order to calculate the RMSE for a longer horizon, we analyzed our result by sampling trajectories directly from the actual motion model during the whole prediction period and compared them with the real trajectory.

For a selected test case, the ground-truth motion model is determined by the final insertion area and the RMSE can be calculates as:

$$RMSE = \sqrt{\frac{1}{N_s \cdot T \cdot f_s} \sum_{i=1}^{N_s} \sum_{j=1}^{T \cdot f_s} (s_j - \hat{s}_j^i)}, \quad (10)$$

where s_j is the true state value at time j and $\hat{s}_j^i$ is the predicted value from the i-th sampled trajectory at the same time step j. We evaluated the RMSE of the lateral position and the longitudinal velocity separately for each motion model. Moreover, we selected the *pass* model to evaluate the performance of using and not using the time-to-lane-change information obtained from the upper module. The RMSE results across different motion models under various prediction horizons are presented in Table I

As shown in Table I, our method has lateral position errors within 0.5m for each model and a 1.3m/s error in velocity when prediction horizon is 4s. For the lateral position error, it can be clearly seen that the performance of the reference vehicle is better than that of the predicted vehicle. This is because during the lane changing period, the predicted car has larger lateral displacement than the reference car which usually remains its lateral position. When the *pass* model does not use the time information given by the SIMP method, the RMSE values increase especially for velocity errors. This happened because without the time constraint, the model may assume both vehicles will remain their speed and none of them will yield to the other at the beginning stage, which

could cause cumulated speed deviations for both vehicles as well as a delay in the lateral displacement for the predict vehicle (yellow). Note that only a 4s prediction horizon is examined for the *keep* model since there won't have much motion changes within that period for car following cases.

B. Visualization of Selected Cases

We selected two distinct traffic situations to visualize our results as shown in Fig. 4. Each situation has a maximum prediction horizon of 4s (40 frames) and four representative frames are selected from each scenario to illustrate the performance of our framework. In order to better visualize the results, we did not add the trajectories generated by the *keep* motion model on the plot.

At the early stage of the situation in Fig.4 (a), the predicted vehicle is likely to interact with its left reference car since the vehicle is closer to the area behind its left car than the area in front of its right car. Therefore, in the first frame, most trajectories are sampled from the *yield* motion model and the yielding pattern can be clearly recognized from the plot. As time goes, the predicted vehicle starts to increase its speed and finally passes its right reference vehicle. For case in Fig.4 (b), the predicted vehicle changes its mind by first having a large intention to pass its left vehicle and then choosing to yield to its right reference vehicle. Note that both cases have some intermediate stages where multiple possible interaction exist and thus it is necessary to predict all these potential joint trajectories to have a full understanding of the scene.

C. Framework Robustness Test

Although our framework shows great performances during the evaluation, it is possible that some corner cases are not included in the test data. Therefore, to examine whether our method can generate reasonable results under various situations, we manually changed some parameter settings to

Model	Vehicle	Prediction horizon					
		4s	3s	2s	1.5s	1s	0.5s
Keep	$pred$	0.30	-	-	-	-	-
	ref	0.28	-	-	-	-	-
Yield	$pred$	0.51	0.48	0.39	0.21	0.14	0.08
	ref	0.42	0.36	0.25	0.18	0.11	0.07
Pass	$pred$	0.43	0.41	0.37	0.21	0.19	0.07
	ref	0.41	0.33	0.22	0.14	0.04	0.02
Pass (w/o time)	$pred$	1.10	0.77	0.67	0.51	0.28	0.12
	ref	0.34	0.30	0.29	0.18	0.07	0.03

(a) RMSE for Lateral Position (m)

Model	Vehicle	Prediction horizon					
		4s	3s	2s	1.5s	1s	0.5s
Keep	$pred$	1.11	-	-	-	-	-
	ref	1.07	-	-	-	-	-
Yield	$pred$	1.24	1.19	0.91	0.79	0.70	0.31
	ref	1.25	1.05	0.72	0.67	0.33	0.22
Pass	$pred$	1.23	1.08	0.89	0.74	0.50	0.25
	ref	1.34	1.01	0.74	0.58	0.49	0.20
Pass (w/o time)	$pred$	2.59	2.25	1.97	1.41	0.90	0.79
	ref	2.17	2.06	1.32	1.08	0.64	0.43

(b) RMSE for Longitudinal Velocity (m/s)

have an comprehensive analysis. We used two methods to evaluate the framework robustness:

- Fix the current input scenario and change motion models in the lower module.
- Fix the motion model in the lower module and change the time-to-lane-change (TTLC) distribution obtained from the upper module.

For the first evaluation, we chose the 10th frame of the test scenario in Fig. 4 (a), where the predicted car has a similar longitudinal position as its right reference car and the actual motion model is *pass*. We then changed the motion model to *yield* and the comparison results of using different models are shown in Fig. 5. As can be seen in the plot, the *pass* model predicts that the predicted vehicle will first increase the speed to surpass the reference vehicle and then perform large lateral deviations for lane change. The *yield* model, however, indicates that the predicted vehicle will decrease its speed at the beginning and yield the reference car while changing its lane. Also, as the model is switched from *pass* to *yield*, the longitudinal velocity of the reference vehicle increases since the predicted car is assumed not to insert in front of it due to the decreasing speed of the predicted car.

We used the same test scenario for the second evaluation but evaluated from the starting frame. By raising the predicted TTLC value from 1s to 4s, the predicted vehicle will have more time to change its lane. According to Fig. 6, when TTLC is 1s, the predicted car has a 0.6m lateral

deviation within one second; when TTLC increases, the vehicle gradually changes its lane, which takes up to 2s to have the same amount of lateral deviation. The speed comparison also reasonably indicates that the vehicle will have a larger acceleration when it decides to change the lane faster.
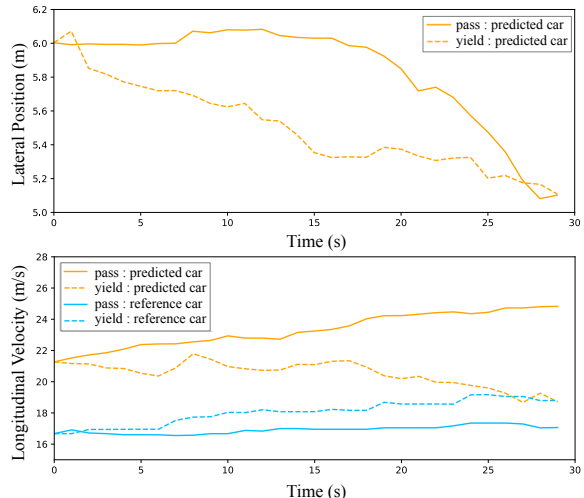

Fig. 5. Performance Comparison using Different Motion Models

Fig. 6. Performance Comparison using Different Temporal Information

V. CONCLUSIONS

In this paper, a framework for probabilistic traffic scene prediction was proposed. It can simultaneously predict possible motions for multiple interacting traffic participants under various circumstances. An exemplar highway scenario with real-world data was used to demonstrate the performance of the framework. The proposed framework is not only able to generate accurate trajectories sampled from the predicted joint distributions of scene entities, but also has the adaptability to different driving situations. Note that we used the

CVAE method to analyze the performance of the framework and it does not necessarily conclude that this is the best suited to the lower module. Other approaches that can generate conditional joint distributions are also able to use and we will compare different methods as well as examine the framework on other scenarios in our future works.

REFERENCES

[1] H. M. Mandalia and D. D. Salvucci, "Using support vector machines for lane change detection," in *Proc. of the Human Factors and Ergonomics Society 49th Annual Meeting*, 2015.

[2] J. C. McCall, D. P. Wipf, M. M. Trivedi, and B. D. Rao, "Lane change intent analysis using robust operators and sparse bayesian learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 3, pp. 431-440, 2007.

[3] J. Li, W. Zhan, and M. Tomizuka, "Generic Probabilistic Interactive Situation Recognition and Prediction: From Virtual to Real," in *2018 IEEE Intelligent Transportation Systems (ITSC)*, to appear.

[4] S. Yoon and D. Kum, "The multilayer perceptron approach to lateral motion prediction of surrounding vehicles for autonomous vehicles," in *2016 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2016, pp. 1307-1312.

[5] M. Liebner, M. Baumann, F. Klanner and C. Stiller, "Driver intent inference at urban intersections using the intelligent driver model", in *2012 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2012, pp. 1162-1167.

[6] R. Izquierdo, I. Parra, J. Muoz-Bulnes, D. Fernndez-Llorca, and M. A. Sotelo, "Vehicle trajectory and lane change prediction using ANN and SVM classifiers," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1-6.

[7] J. Wiest, M. Höffken, and U. Kreßel, and K. Dietmayer, "Probabilistic Trajectory Prediction with Gaussian Mixture Models," in *2012 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2015, pp 141-146.

[8] F. Altché, and A. De La Fortelle, "An LSTM network for highway trajectory prediction," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC): Workshop*. IEEE, 2017.

[9] D. Lenz, F. Diehl, M. T. Le, and A. Knoll, "Deep neural networks for Markovian interactive scene prediction in highway scenarios," in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 685-692.

[10] B. Kim et al., "Probabilistic Vehicle Trajectory Prediction over Occupancy Grid Map via Recurrent Neural Network," arXiv preprint arXiv:1704.07049, 2017.

[11] S. Park, B. Kim, C. M. Kang, C. C. Chung, and J. W. Choi, "Sequence-to-Sequence Prediction of Vehicle Trajectory via LSTM Encoder-Decoder Architecture," in arXiv:1802.06338, 2018.

[12] S. Klingelschmitt, F. Damerow, V. Willert, and J. Eggert, "Probabilistic situation assessment framework for multiple, interacting traffic participants in generic traffic scenes," in *2016 Intelligent Vehicles Symposium (IV)*, 2016, pp. 1141-1148.

[13] S. Klingelschmitt, V. Willert, and J. Eggert, "Probabilistic, discriminative maneuver estimation in generic traffic scenes using pairwise probability coupling," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1269-1276.

[14] Y. Hu, W. Zhan, and M. Tomizuka, "Probabilistic prediction of vehicle semantic intention and motion," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, pp.307-313.

[15] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," In *Advances in Neural Information Processing Systems (NIPS)*, pp 3483-3491, 2015

[16] J. Walker, C. Doersch, A. Gupta, and M. Hebert, "An uncertain future: Forecasting from static images using variational autoencoders," in *ECCV*, 2016

[17] U.S. Department of Transportation Intelligent Transportation Systems Joint Program Office (JPO). Avaliable: https://www.its.dot.gov/data/