Impact of Traffic Lights on Trajectory Forecasting of Human-driven Vehicles Near Signalized Intersections

Geunseob (GS) Oh, Huei Peng

Abstract— Forecasting trajectories of human-driven vehicles is a crucial problem in autonomous driving. Trajectory forecasting in the urban area is particularly hard due to complex interactions with cars and pedestrians, and traffic lights (TLs). Unlike the former that has been widely studied, the impact of TLs on the trajectory prediction has been rarely discussed. Our contribution is twofold. First, we identify the potential impact qualitatively and quantitatively. Second, we present a novel resolution that is mindful of the impact, inspired by the fact that human drives differently depending on signal phase and timing. Central to the proposed approach is *Human Policy Models* which model how drivers react to various states of TLs by mapping a sequence of states of vehicles and TLs to a subsequent action of the vehicle. We then combine the Human Policy Models with a known transition function (system dynamics) to conduct a sequential prediction; thus our approach is viewed as *Behavior Cloning*. One novelty of our approach is the use of vehicle-to-infrastructure communications to obtain the future states of TLs. We demonstrate the impact of TL and the proposed approach using an ablation study for longitudinal trajectory forecasting tasks on real-world driving data recorded near a signalized intersection. Finally, we propose *probabilistic* (generative) Human Policy Models which provide probabilistic contexts and capture competing policies, e.g., *pass* or *stop* in the yellow-light dilemma zone.

I. INTRODUCTION

Autonomous driving has been more successful in highway than in urban city mainly due to the simplicity of its driving environment; absence of traffic signals, and more stable interactions with other vehicles. Realizing fully autonomous vehicles in urban driving environments is more challenging for the opposite reasons.

One of the major differences between urban city and highway driving is traffic lights (TLs). In urban areas, especially in the vicinity of TLs exemplified by signalized corridors or intersections, the motions of vehicles are mainly governed by traffic signals. People obey the traffic signals and properly respond to implicit rules imposed by traffic lights. Examples of the implicit traffic rules include stopping for a traffic light in a red phase, maintaining a proper speed in a green phase in a free-flow situation. This is why predicting how human drivers respond to traffic signals is the key to the trajectory forecasting in urban area. The decision-making, path planning, and control synthesis all benefit from more accurate trajectory forecasting, ultimately leading to successful autonomous driving.

Recent studies in trajectory forecasting utilize generative models (e.g., variational autoencoders (VAE) [1] or gener-

Geunseob (GS) Oh and Huei Peng are with the Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109 USA, Email : {gsoh, hpeng}@umich.edu
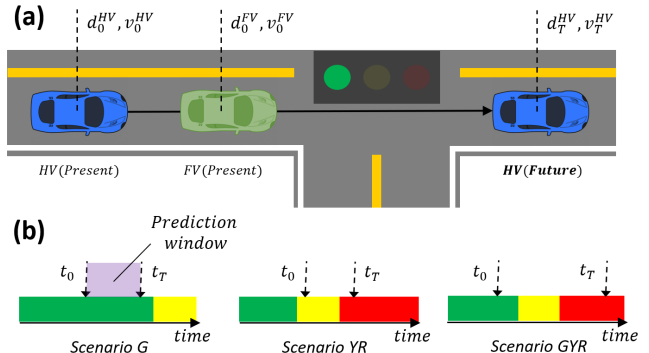
Fig. 1. (a) The trajectory forecasting problem near a traffic light for the vehicles with *through* moves is depicted. Given a sequence of past states of a host vehicle (HV) and contexts (states of TLs, its front vehicle (FV)), our goal is to forecast HV's states under various states of TLs. (b) Three example scenarios of the problem are described. The full list of the scenarios is described in Table. I. We define 'scenario G' as a forecasting problem when the prediction window starts on a green light and ends on the same green light. 'Scenario GYR' represents a forecasting problem where the window spans over a set of green, yellow, and red lights.

ative adversarial networks (GAN) [2], [3]) or convolutional & recurrent neural-nets models [4], [5], [6], [7], [8], [9]. The existing works mainly focus on accurate modeling of the interactions among vehicles [1], [4], [5], [6], [10], and/or pedestrians [2], [7], [8], [9]. Despite the important role of TLs in the vehicle motions, efforts to understand the dynamics between TLs and vehicle trajectories and to quantify the impacts of TLs on the trajectories have barely been made in the trajectory forecasting literature.

The less recognized impact of TLs in the vehicle trajectory forecasting is elaborated in Fig. 2. Specifically, Fig. 2 depicts an example which shows how the states of TLs affect vehicle trajectories and how uncertainties in the states of TLs can cause high prediction errors. Even for models that account for the uncertainties in the future phase (it can either be red or green) and output probabilistic predictions for the two possibilities (either the phase remains red, or the phase shifts to green), the uncertainty question still is not resolved: precisely when will the phase change?

On the other hand, there have been efforts to model the dynamic impact of TLs in the transportation research community. However, there is no comprehensive model that describes behavior of human drivers near traffic signals. A few papers have studied specific instances of the dynamics but limited to a few simple scenarios; [11], [12], [13] developed models for vehicles approaching a signalized intersection (SI) and making complete stops in red light. [13], [14], [15], [16], [17] proposed models for vehicles departing
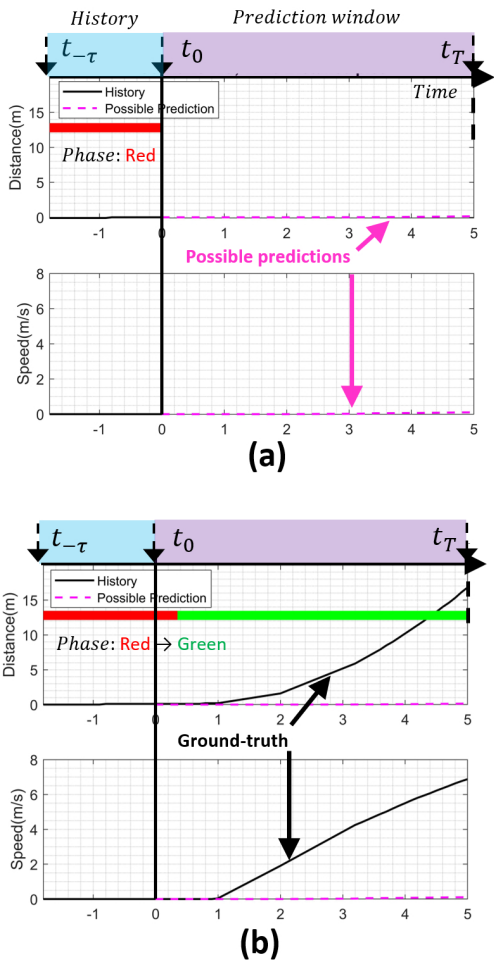
Fig. 2. A motivational example from our real-world dataset is depicted to demonstrate how uncertainty in the states of TLs makes the trajectory forecasting difficult. The sample is a 7s long RG scenario, where the first 2s is the observed history and the last 5s is the future. The goal is to predict *longitudinal* position and speed of a vehicle for $t = 0s : 5s$. We define $0m$ as the position of the vehicle at $t = 0s$. The distance to TL is then $13m$. (a) Given the history of the vehicle's position, speed, and the red phase observed at $t = 0s$, a reasonable prediction under uncertain traffic phase at $t > 0s$ that an existing method would make is to predict the vehicle to stay put (pink dotted line). (b) However, in reality, the phase shifted to green at $t = 0.35s$ and the vehicle sped away as depicted as black line, resulting in large prediction errors in both position and speed.

from a SI from zero-speed in green phase. These models are either limited to specific instances of the problem, or are not forecasting algorithms since they require parameters like total deceleration time, final speed, which can only be measured after a trip is complete. [18], [19] presented prediction algorithms for the vehicles in highway and in car-following scenarios [20] based on car models proposed in [21], [22]. These models, however, do not describe how human drivers react to the traffic signals.

In order to leverage rich information that comes from the dynamics between TLs and human drivers and to apply it to the trajectory forecasting problem, we simply utilize vehicle-to-infrastructure (V2I) communications. By leveraging V2I communication, one can access **the future profile** of the TLs ahead of time, which resolves the aforementioned uncertainty problem inherent to any prediction task near traffic lights.

Based on this simple but novel idea, we investigate how the access to the future profile of the TLs can improve the accuracy of the trajectory forecasts.

In this paper, we tackle a trajectory forecasting problem described in Fig. 1. In order to solely focus on the impact of TLs, the forecasting task is simplified to a longitudinal trajectory forecasting with a front vehicle (FV) near a TL. This setting assumes that the impact of rear & side vehicles is minimal and no vehicle cuts into the lane that HV is located.

Our solution approach consists of two models which map a sequence of states of a HV, its FV and the corresponding states of a TL to the subsequent action (longitudinal acceleration) that the HV takes. We name the models *Human Policy Models* to highlight that the models return an action given a state. Since the models are learned to mimic drivers, the proposed approach can be viewed as *imitation learning*, or precisely *behavior cloning*, as we work with real data (i.e., no access to environment) and assume that inputs are i.i.d.

The first model is *deterministic* human policy model which returns the most-probable action and is designed using RNN-based network. The second model is *probabilistic* (generative) human policy model which outputs distribution parameters and is able to generate trajectories by sampling from the distribution and is designed using a RNN-based mixture density network (MDN) [23]. Then, we utilize these models to forecast longitudinal trajectories of the HV sequentially over a time span using the system dynamics (trasition function). For the training, validation, and testing, we used 502,253 sequences excerpted from naturalistic (non-obstructive, uncontrolled) trips from 50 distinct cars over 2 years at a signalized intersection in Ann Arbor, MI.

The remainder of the paper is organized as follows: Section II elaborates on the proposed human policy models. Section III describes the framework is used to obtain the forecast trajectories. Section IV presents results using an ablation study. Finally, Section V offers concluding remarks.

II. HUMAN POLICY MODEL

A. Problem Description

The problem of longitudinal trajectory forecasting under various phases and timings of TLs (Fig. 1) is challenging due to the stochastic reactions of human drivers. For example, a driver may prefer late hard-braking approaching to a red light, while others may prefer soft-braking. Also, the reactions of drivers at steady phases (i.e., green (G), yellow (Y), red (R)) are different from those at phase transitions (GY, YR, RG). Another example is the decision making in yellow light dilemma zone [24], where a driver arrives at a TL at a high speed. There usually exists two competing decisions; a driver could either engage in a hard-braking to stop before the TL or pass through the TL.

In this sense, we break down the problem to six distinct scenarios depicted in Fig. 1(b) and Table. I. The idea behind this categorization is our belief that humans react differently to various phases of TLs, resulting in unique trajectories. In this regard, we argue that a *comprehensive* model should be validated against all the 6 scenarios.

B. Related Works

To the best of our knowledge, none of the existing works in machine learning community addressed this particular forecasting problem (see Section I). In traditional transportation community, a few papers have discussed acceleration models or speed profiles near traffic signals. [13], [15], [16], [17] proposed polynomial speed & acceleration models for vehicles departing from a SI from zero-speed in G phase. [11], [12], [13] developed deceleration models for vehicles make complete stops at a TL in R phase. However, these models studied very specific instances of the problem, thus do not qualify as comprehensive model. We classified the available studies in Table. I.

TABLE I

Six distinct scenarios of the prediction problem

	Scenario	Available Studies in Transportation Community
G	D0 (departure from zero-speed)	ATL Newzealand(1990), Bham(2002), Day(2013), Modified IDM(2018)
	General	None
Y		None
R	A0 (arrival to zero-speed)	Bennett(1995), Wang(2005), Modified IDM(2018)
	General	None
GY		None
YR		None
RG		None

We believe that a general model which captures the behavior of human drivers in all scenarios described in Table. I is crucial to accurate forecasting of human vehicles near TLs. To the best of our knowledge, such model does not exist.

C. Proposed Model

The key behind modeling a driver's reaction to TL is feature selection and model design so that the reaction to TL is well captured in corresponding state space. We first introduce the model, define the state, and move on to the explanation of the intuition behind the selections.

Human Policy Models are functions that map a sequence of past states of a HV ($X_{t-\tau:t}^{HV} := [d_{t-\tau:t}, v_{t-\tau:t}]$) and a context vector ($C_t := [X_t^{FV}, X_t^{TL}, TOD_t]$) to a subsequent longitudinal acceleration of HV (a_t^{HV}). d_t and v_t each indicates the distance to the traffic light and speed at time t. C_t includes state of the FV at time t ($X_t^{FV} := [FV_t, r_t, \dot{r}_t]$) where FV_t is a binary flag which represents the presence of a FV, and $r_t, \dot{r}_t$ indicate positions and speed relative to the HV. The state of the corresponding TL at time t ($X_t^{TL} := [P_t, T_t]$) is defined as phase (G,Y,R) as P_t and timing (i.e., time elapsed in the current phase) of TLs as T_t. TOD_t is the time of day at time t. The intuitions behind the selection of the input features ($X_t := [X_{t-\tau:t}^{HV}, X_t^{FV}, X_t^{TL}, TOD_t]$) are as follows.

Distance to traffic light (d_t), **Speed** (v_t) each represents the longitudinal distance of a HV to the TL that the HV is

approaching or departing from and the longitudinal speed of the HV. They are essential in forecasting vehicle trajectories near TLs. For example, a HV approaching a TL in red phase travels slowly when it is close to the TL, whereas it can travel fast when it is far away from the TL. $d_x > 0$ means that the HV is approaching the TL, and $d_x < 0$ indicates that it's departing from the TL. We assume $v_x >= 0$.

Range and range-rate ($r_t, \dot{r}_t$) represent the longitudinal position and speed of the FV relative to those of the HV. We assume $r_x > 0$, meaning that the FV is always ahead of the HV. Note, rear & side vehicles are not considered due to the dataset availability. However, one can trivially extend our model to include them.

Phase and timing of traffic light (P_t, T_t) represents the phase of a TL (G,Y,R) that a HV is subject to and the time elapsed since the last phase change ($T_t >= 0$). T_t accounts for transient behaviors of human drivers at phase shifts. For example, a vehicle approaching a TL in a red phase with a small T_t, meaning that the phase just shifted to red, may not be traveling slowly whereas a vehicle approaching a TL with a large T_t is likely to travel slowly or fully stopped.

Time of day (TOD) represents the time of day as elapsed hours since midnight ($0 \leq TOD < 24$). TOD=12 indicates noon. Macroscopic traffic characteristics including traffic speed differ considerably depending on TOD as evidenced in studies including [25]. The selection of TOD is an attempt to incorporate the macroscopic trend of the traffic.

Due to stochastic and complex nature of human decision making in driving, a simple analytical model such as a polynomial or a physics-based model may not represent the nominal or probabilistic behaviors of human-drivers near traffic signals well. This is why we opt for data-driven approach.

D. Dataset

The real-world driving data utilized in this work are from Safety Pilot Model Deployment (SPMD), a large-scale connected vehicle study conducted in Ann Arbor, MI [26]. The vehicles were equipped with data loggers which collected 10Hz GPS signals including coordinates, speeds, and heading angles as well as 10Hz front vehicles data such as relative positions and speeds. While SPMD database does not include any vision data (e.g., lidar, camera, or radar) nor the access to all vehicles in the scene, it provides detailed TL profiles that were obtained using V2I communication devices installed at signalized intersections (SIs) and information about vehicles in front of HV. To the best of our knowledge, no dataset is publicly available that provides the detailed TL data. Hence, we leverage SPMD for its unique access to the TL profiles.

In this work, we extracted 502,253 observations (samples) from 50 distinct SPMD vehicles collected over a span of 27 months (Mar 2015 - July 2017) near the Plymouth Rd & Huron Pkwy intersection. Each observation was synchronized with the traffic signal states of the SI. In order to reduce the noise in $X_t^{HV}, X_t^{FV}, a_t^{HV}$, a least-square polynomial smoothing filter was used [27].
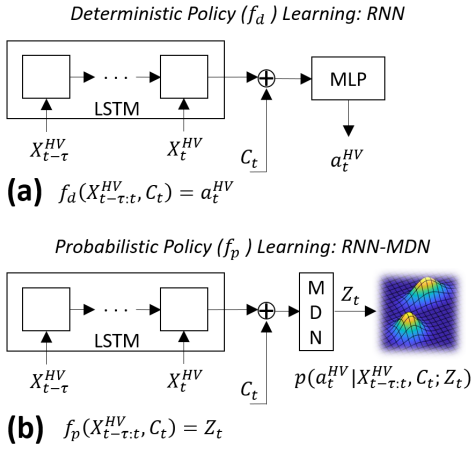
(a) $f_d(X_{t-\tau:t}^{HV}, C_t) = a_t^{HV}$

(b) $f_p(X_{t-\tau:t}^{HV}, C_t) = Z_t$

Fig. 3. The proposed policy models are described. (a) a deterministic policy ($f_d : [X_{t-\tau:t}^{HV}, C_t] \to a_t^{HV}$), (b) a probabilistic policy ($f_p : [X_{t-\tau:t}^{HV}, C_t] \to Z_t$).

E. Implementation Details

Both the deterministic (f_d) and probabilistic (f_p) human policy models were implemented in Tensorflow-Keras. Each model consists of a double-stacked LSTM which takes $X_{t-\tau:t}^{HV}$ followed by a concatenation with the context vector $C_t = [X_t^{FV}, X_t^{TL}, TOD_t]$. The concatenated tensor is then fed into a multi-layer perceptron (MLP) for the deterministic model or a MDN for the probabilistic model. The MLP layer outputs a_t^{HV} whereas the MDN layer outputs the distribution parameters Z_t. As we model MDN using gaussian mixtures, the MDN layer outputs three sets of the parameters: mixture weights π_k, means μ_k, variances σ_k^2 for N components. We used $N = 2$ for the yellow light dilemma scenario. Models were trained using ADAM optimizer.

f_d is learned by minimizing a loss function L_d which is a summation of mean squared error as described below.

$$L_d := \sum_{t=1}^{T}(a_t^{HV} - f_d(X_{t-\tau:t}^{HV}, C_t))^2 \quad (1)$$

f_p is obtained by minimizing a loss function L_p which is a sum of a negative log-likelihood.

$$L_p := \sum_{t=1}^{T} -log(p(a_t^{HV}|X_{t-\tau:t}^{HV}, C_t; Z_t)) \quad (2)$$

III. TRAJECTORY FORECAST FRAMEWORK

Fig. 4 illustrates the proposed framework that consists of two parts. The first part is an off-line supervised learning of the policies. The second part is where we use the learned policies in sequential prediction setting to obtain $X_{0:T}^{HV}$ (i.e., alternation of one-step predictions of the learned policies and state transitions). The system dynamics (i.e., transition function for X^{HV}) is given as longitudinal vehicle kinematics with a zero-order hold, as described in Eq. 3 and 4.

$$v_{n+1} := v_n + a_n\Delta t_n, d_{n+1} := d_n + 0.5(v_{n+1} + v_n)\Delta t_n \quad (3)$$

$$X_{n+1}^{HV} = A_n X_n^{HV} + B_n a_n^{HV} \quad (4)$$

where $A_n := \begin{bmatrix} 1 & \Delta t_n \\ 0 & 1 \end{bmatrix}, B_n := \begin{bmatrix} 0.5\Delta t_n \\ \Delta t_n \end{bmatrix}, X_n^{HV} = \begin{bmatrix} d_n \\ v_n \end{bmatrix}.$
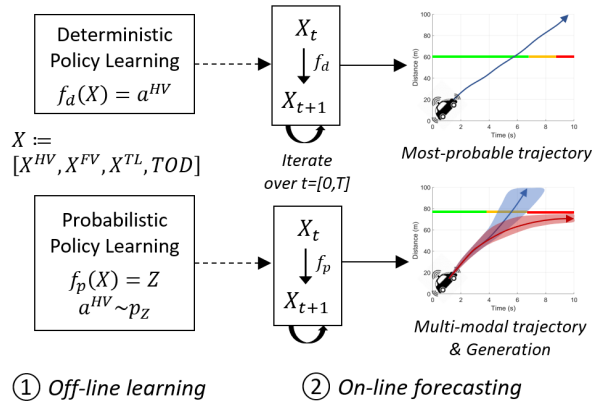

Fig. 4. The framework is divided into two steps. The first step is the off-line training of the proposed models. The second step involves propagations of the policies and state transitions to obtain vehicle trajectories over T.

Defining $n = 0$ and $n = N$ as the index for $t = 0$ and $t = T$, the trajectory forecast over the prediction horizon $t = [0, T]$ are obtained by propagating Eq. 4 from $n = 0$ to $n = N - 1$:

$$
\begin{aligned}
X_N^{HV} &= \prod_{k=0}^{N-1} A_k X_0^{HV} + \prod_{k=1}^{N-1} A_k B_0 a_0^{HV} + \prod_{k=2}^{N-1} A_k B_1 a_1^{HV} \\
&\quad + ... + A_{N-1}B_{N-2}a_{N-2}^{HV} + B_{N-1}a_{N-1}^{HV} \\
&= \prod_{k=0}^{N-1} A_k X_0^{HV} + \prod_{k=1}^{N-1} A_k B_0 f(X_{-n_\tau+1:0}^{HV}, C_0) \\
&\quad + ... + B_{N-1}f(X_{N-n_\tau:N-1}^{HV}, C_{N-1}) \\
&:= F(X_{1:N-1}^{HV}, X_{-n_\tau+1:0}^{HV}, C_{0:N-1}, \Delta t_{0:N-1})
\end{aligned}
\quad (5)
$$

where f can either be f_d or $S(f_p)$. S is a function which returns a sample a_t^{HV} from the pdf. In case of 1D Gaussian, $Z_t := [\mu_t, \sigma_t^2]$ and $a_t^{HV} \sim N(\mu_t, \sigma_t^2)$. τ, n_τ each indicates input sequence length in time and in the number of steps.

As described in Eq. 5, X_N^{HV} is a function (F) of $[X_{1:N-1}^{HV}, X_{-n_\tau+1:0}^{HV}, C_{0:N-1}, \Delta t_{0:N-1}]$. The second term $X_{-n_\tau+1:0}^{HV}$ is given and the last term $\Delta t_{0:N-1}$ can simply be predetermined based on a required time resolution. Obtaining $C_{0:N-1}$ at the prediction time ($t = 0$) is the main challenge, due to uncertainties in $X_{1:N-1}^{FV}, X_{1:N-1}^{TL}$. A simple way to get away with the uncertainties is to design a model to predict trajectories ($X_{0:T}^{HV}$) conditioned only on the observed states $[X_{-\tau:0}^{FV}, X_{-\tau:0}^{TL}]$. An example is a model with many-to-many RNN that takes a sequence of past states and returns a sequence of future states; which is a forecasting model that does not utilize the future states of TLs. Fig. 2 showed how such model can fail.

This is where our novel idea comes into play. We remove the uncertainties by utilizing the future phases and timings of TLs obtained via V2I communications. With the access, $X_{1:N-1}^{TL}$ can be attained at the prediction time. The remainder is then $X_{1:N-1}^{FV}$, which is obtained using a variant of f_d. Specifically, we train another human policy model f_d^{NoFV}: $[X_{N-n_\tau:N}^{HV}, C_N'] \to a_N^{HV}$ with $C_N' := [X_N^{TL}, TOD_N]$ excluding X_N^{FV} from C (i.e., f_d^{NoFV} does not consider FV). After the off-line learning of f_d^{NoFV}, we apply the aforementioned iterative process on FV to obtain $X_{1:N-1}^{FV}$ via Eq. 5 with f_d^{NoFV}.
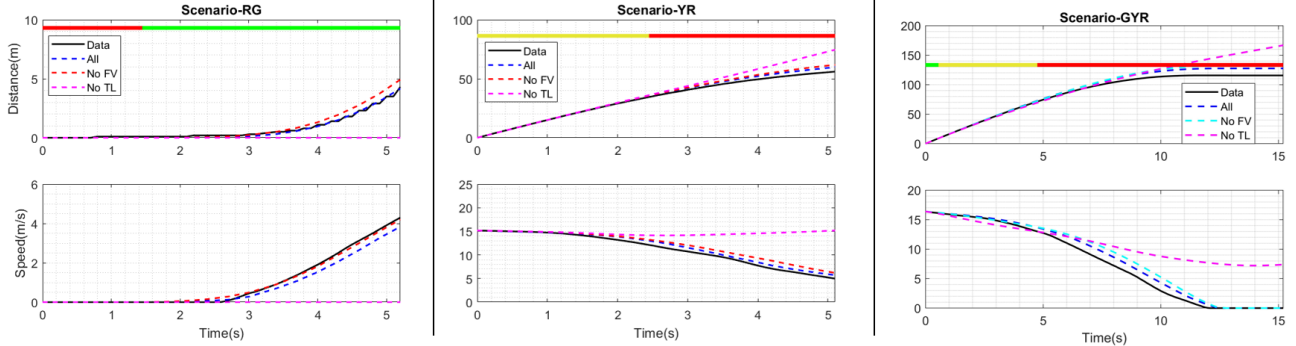
Fig. 5. For qualitative evaluations, ground-truth and predicted trajectories are depicted for two 5s sample scenarios RG (left), YR (middle), and one 15s sample scenario GYR (right). Four trajectories of distance to the TL (1st row), speed (2nd row) are shown in each plot and represent ground-truth (black), predictions $X_{0:5s}^{HV}$ obtained using f_d(Blue), f_d^{NoFV}(Red), f_d^{NoTL}(Pink). The 2s history $X_{-2:0s}^{HV}$ (i.e., part of inputs to the models) are omitted for the simplicity.

Once $X_{1:N-1}^{FV}, X_{1:N-1}^{TL}$ are secured, the resulting trajectory forecast $X_{1:N}^{HV}$ is obtained via Eq. 5. Since $X_{1:N}^{HV}$ can simply be forecast using f_d^{NoFV}, we conduct an ablation study (elaborated in Section IV) on f_d, f_d^{NoFV} and the other two models (f_d^{NoTL}, f_d^{NoFVTL}) which each represents unique model where X^{TL} and $[X^{FV}, X^{TL}]$ are excluded from C.

For the probabilistic human policy model, the probability of the resulting trajectory forecast $p(X_{1:N}^{HV})$ can be estimated using the chain rule of probability, which factorizes the joint distribution over N separate conditional probabilities:

$$p(X_{1:N}^{HV}|X_{-n_\tau+1:0}^{HV}, C_{0:N-1}, \Delta t_{0:N-1}) = \prod_{k=1}^{N} p(X_k^{HV}|X_{1:k-1}^{HV}, X_{-n_\tau+1:0}^{HV}, C_{0:k-1}, \Delta t_{0:k-1}) \quad (6)$$

As opposed to the deterministic forecasting where the most-probable trajectory is obtained, a resulting trajectory is a sample from a probability distribution. While we can estimate the joint probability density of a trajectory forecast via Eq. 6, the marginal probability of X_t^{HV} needs to be numerically estimated via sampling since the distribution parameter Z_t is obtained via an arbitrarily complex neural network and depends on previous predictions $X_{-n_\tau+1:k-1}^{HV}$. Thus, we utilize Monte Carlo Simulation to obtain the samples (roll-out trajectories) and kernel density estimation to approximate the marginal probability density of the samples.

IV. RESULTS

In this section, we discuss evaluation results conducted on the dataset (see Section II-D). In Section IV-A, we present the impact of X^{TL} qualitatively by comparing the trajectory forecasts made using the four deterministic policies and that the utilization of phases and timings of TLs helps forecasting trajectories more accurately. In Section IV-B, we discuss a set of metrics to evaluate performance of the forecasts. The ablation study in Section IV-C is designed to evaluate the 4 variants of our deterministic models ($f_d, f_d^{NoFV}, f_d^{NoTL}, f_d^{NoFVTL}$) by quantifying their performance using 3,111 test snippets. In Section IV-D, we demonstrate how the proposed probabilistic models can tackle scenarios with competing policies.

A. Impact of X^{TL} on trajectory forecasts

As explained in Section III, we first train 4 distinct deterministic policy models $f_d, f_d^{NoFV}, f_d^{NoTL}$, and f_d^{NoFVTL}. The superscript specifies the context input C that the model takes (i.e., $C^{NoFVTL} := [TOD], C^{NoTL} := [X^{FV}, TOD], C^{NoFV} := [X^{TL}, TOD]$, where C^{mode} is the context input for f_d^{mode}). Each model then produces unique trajectory forecasts per scenario. Three scenarios (RG, YR, GYR) are sampled and the four models are used to forecast trajectories. In each sample, four trajectories are depicted to represent ground-truth and 3 predicted trajectories each from $f_d, f_d^{NoFV}, f_d^{NoTL}$.

The sample on the left of Fig. 5 is similar to the motivational example in Fig. 2. At $t = 0$, the driver was at a stop at a red phase. A reasonable prediction that a model without $X_{0:5s}^{TL}$ would make is to forecast the vehicle to stay put. As expected, f_d^{NoTL} (pink) failed to predict X^{HV} accurately, whereas the forecasts from the other two models f_d, f_d^{NoFV} which utilize $X_{0:5s}^{TL}$ are close to the ground-truth.

The middle plot of Fig. 5 is a scenario YR where the driver was cruising as approaching the TL ($v_{-2:0s}^{HV} = 15$). Given $P_0 = Y$ and that the vehicle was cruising, f_d^{NoTL} forecast the vehicle to maintain the speed, causing the prediction errors to grow over time. The other two models f_d, f_d^{NoFV} which use $X_{0:5s}^{TL}$ took account for the phase shift in the future and accurately forecast how the driver would react to the shift.

On the other hand, the right plot serves as an exemplar long-term scenario. A 15s scenario that spans a full cycle of phases (GYR) is illustrated where HV was initially decelerating while approaching the TL. All models captured temporal trends in the speed and predicted that HV would continue to decelerate. f_d^{NoTL} predicted that the vehicle would cross the intersection, given $P_0 = G$ whereas f_d, f_d^{NoFV} predicted that the vehicle would make a stop before the TL considering the future states of TL. Indeed, the ground-truth is that the vehicle made a stop before the TL to react to the phase shift.

As demonstrated, the impact of X^{TL} is significant; uncertainties in X^{TL} can cause high prediction errors, especially for long-term predictions. The results suggest that forecasting methods may perform poor without knowledge of future X^{TL}, highlighting why the problem is critical. While information
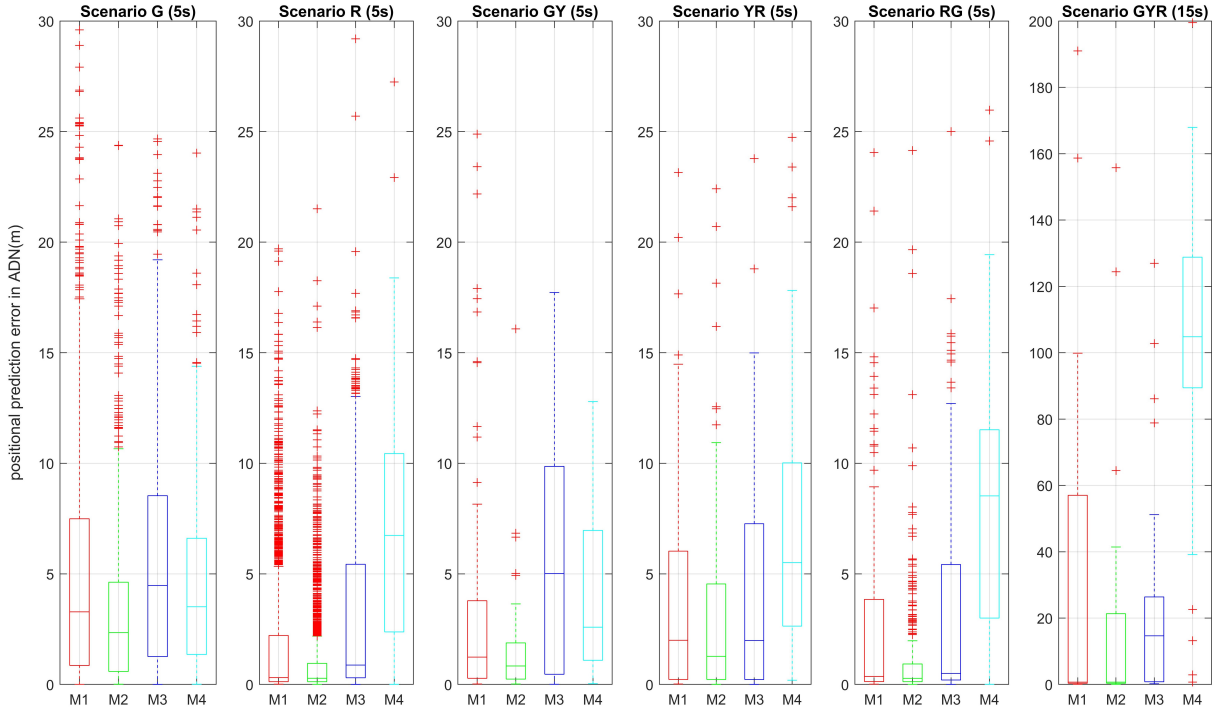
Fig. 6. A quantitative evaluation is conducted on the test set and depicted as boxplots for the six scenarios using the evaluation metric *ADN*. M1, M2, M3, M4 each represents forecasting models with $f_d, f_d^{NoFV}, f_d^{NoTL}, f_d^{NoFVTL}$. The prediction window is $T = 15s$ for scenario GYR, and $T = 5s$ for others.

of TLs can be transmitted through FV to some extent, such transmission is ineffective when (1) FV is absent, (2) FV is present, but, far from HV, or (3) FV is present and close to HV, however, the phase transition is imminent. The proposed idea is a solution to the problem: utilization of the future X^{TL} greatly improves the quality of forecasts near traffic lights.

B. Evaluation metrics

We use the following metrics for the quantitative evaluation: mean absolute error (MAE), time weighted absolute error (TWAE), absolute deviation at the end of the prediction window (ADN) defined in Eq. 7, 8, and 9, where $\hat{X}_k^{HV}, X_k^{HV}$ represents the kth-step forecast X and ground truth X.

$$MAE := \frac{\sum_{k=1}^{N} |\hat{X}_k^{HV} - X_k^{HV}|}{N} \tag{7}$$

$$TWAE := \frac{\sum_{k=1}^{N} (t_k |\hat{X}_k^{HV} - X_k^{HV}|)}{\sum_{k=1}^{N} t_k} \tag{8}$$

$$ADN := |\hat{X}_N^{HV} - X_N^{HV}| \tag{9}$$

We used $\forall k : \Delta t_k = 0.2s, \tau = 2s$ (history). For the scenario with a prediction window $t_N = 5s$, the last index N is 25.

C. Ablation Study

The goal of the ablation study is to quantify the impacts of X^{TL} on a large testset and to evaluate performance of the four policies. The result (on ADN) is presented in Fig. 6. The sample sizes are 688 (G), 1909 (R), 68 (GY), 81 (YR), 362 (RG), 32 (GYR), totalling 3,111 sample snippets. The detailed results for the ablation study on all three metrics

(MAE, TWAE, ADN) are presented in Table II, III, each for position and velocity errors. Note, scenario Y is not depicted due to its short (and inconsistent) prediction horizon; we observed that the phase Y usually lasts anywhere between 2.5s to 4s. The first 2s are used as inputs, which means the prediction horizon for the scenario Y is only 0.5s to 2s.

As depicted in Table II, III, the magnitude of error is as follows: MAE<TWAE<ADN, making ADN the largest error. As shown in Fig. 6, across all scenarios, the two models f_d, f_d^{NoFV} which utilizes X^{TL} outperform the other two models f_d^{NoTL}, f_d^{NoFVTL} which don't take advantage of the future phases and timings information. Interestingly, the winner is not f_d, but it is f_d^{NoFV}, which performs the best on all characteristics of boxplot including the 1st, 3rd quartiles, the median, and the upper limit of the extreme points. Our interpretation is that the exclusion of X^{FV} from C increases the prediction accuracy, due to the uncertainty in $\hat{X}_{t>0}^{FV}$ as $\hat{X}_{t>0}^{FV}$ are predicted via another human policy model.

The numbers presented in Table II, III agree with the results from Fig. 6 across all scenarios. f_d^{NoFV} is the winner for almost all metrics, or at least on par with f_d. In summary, the knowledge of future states of traffic lights significantly increase the accuracy of trajectory forecasts, as evidenced in the ablation studies: trajectory forecasts with the winner model have roughly 1.5-30 times smaller (position) MAE, TWAE, ADN for $T = 5s$ scenarios (G, R, GY, YR, RG), and roughly 9-150 times smaller MAE, TWAE, ADN for $T = 15s$ scenario (GYR), compared to trajectory forecasts via f_d^{NoTL}, f_d^{NoFVTL}. This discrepancy in the accuracy becomes bigger as the prediction horizon grows, especially for the long-term forecasts such as scenario GYR.

TABLE II

Ablation study on **position** errors with average MAE, TWAE, ADN. Prediction horizon is 15s for GYR and 5s for others. The lower a metric is the better. The numbers from the best performing model are marked in **bold**.

Scenario	MAE(m)				TWAE(m)				ADN(m)			
	All	NoFV	NoTL	NoFVTL	All	NoFV	NoTL	NoFVTL	All	NoFV	NoTL	NoFVTL
G	1.18	**0.78**	1.35	1.19	1.86	**1.21**	2.16	1.78	3.28	**2.34**	4.47	3.51
R	0.16	**0.15**	0.38	1.79	0.21	**0.20**	0.56	2.98	0.31	**0.28**	0.88	6.74
GY	0.53	**0.41**	1.26	0.82	0.77	**0.54**	2.06	1.31	1.24	**0.84**	5.01	2.59
YR	1.18	**0.69**	0.86	1.01	1.60	**0.93**	1.32	1.42	2.00	**1.27**	1.98	2.29
RG	0.17	**0.14**	0.22	2.28	0.24	**0.19**	0.31	3.80	0.36	**0.28**	0.50	8.53
GYR	**0.445**	**0.445**	4.11	32.40	**0.567**	0.568	6.82	51.16	**0.694**	**0.694**	14.65	104.81

TABLE III

Ablation study on **velocity** errors with average MAE, TWAE, ADN. Prediction horizon is 15s for GYR and 5s for others. The lower the better.

| Scenario	MAE(m/s)				TWAE(m/s)				ADN(m/s)			
	All	NoFV	NoTL	NoFVTL	All	NoFV	NoTL	NoFVTL	All	NoFV	NoTL	NoFVTL
G	0.70	**0.46**	0.97	0.75	1.02	**0.65**	1.41	1.07	1.48	**0.94**	2.28	1.51
R	**0.005**	**0.005**	0.17	1.44	0.007	**0.006**	0.22	2.29	0.010	**0.009**	0.29	4.10
GY	0.39	**0.15**	1.11	0.53	0.54	**0.20**	1.78	0.79	0.87	**0.27**	3.08	1.21
YR	**0.367**	0.374	0.53	1.01	**0.47**	0.52	0.71	1.42	**0.59**	0.69	0.89	2.29
RG	0.007	**0.005**	0.10	1.68	0.010	**0.007**	0.14	2.63	0.02	**0.01**	0.22	4.52
GYR	0.015	**0.014**	1.26	6.78	**0.018**	0.020	1.86	9.80	0.030	**0.029**	2.60	14.13
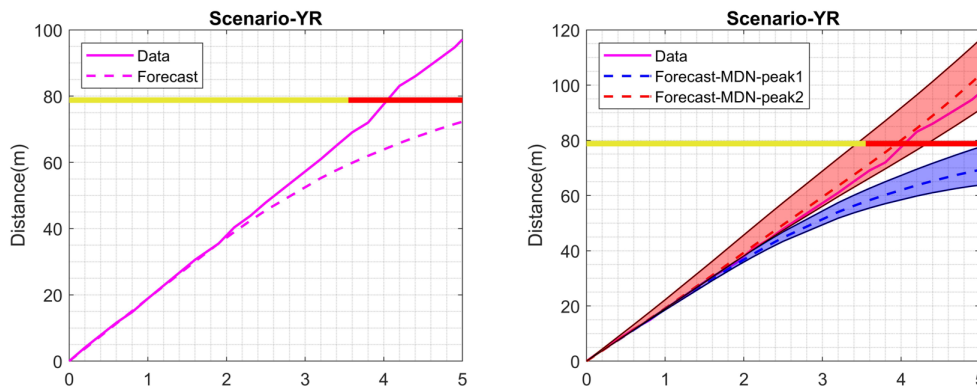
Fig. 7. A sample trip for the yellow light dilemma scenario. The left plot highlights the limitation of the most-probable trajectory. The right plot shows that the proposed probabilistic models are able to forecast the two competing policies. The trajectories $\forall t : p(d_t) >= 0.01$ are illustrated.

D. Probabilistic Prediction

The outliers observed (indicated as '+') in Fig. 6 occur mostly from edge cases and competing policies. Examples of the edge cases include a driver approaching a TL in $P_{\forall t} = R$ with high speed and executing a sudden break right before the TL rather than gradually slowing down as it approaches the intersection. Another example is that a driver in the middle of the road in $P_{\forall t} = G$ moving much slower than the average speed of the traffic for unknown reasons. The outliers occur from competing policies are exemplified by the yellow light dilemma scenario where a driver can either cross the intersection or stop before the intersection.

Fig. 7(a) describes a sample trip observed in our dataset that represents the yellow light dilemma scenario. As shown in the figure, the most-probable trajectory forecast obtained via f_d^{NoFV} predicts that the vehicle would make a stop before the intersection, however, the driver crossed the intersection even after the phase shifted to red.

This is where the proposed probabilistic models come in handy. As the probabilistic model is capable of reproducing multi-modal distributions, it captures the other competing policy (*cross*). In addition, it is able to reason uncertainties of the forecasts as the probabilities of the forecasts can be estimated via Monte-Carlo simulation. Another advantage is that we can generate (sample) possible trajectories as the proposed probabilistic policy outputs both likelihood $p(X|Z)$ and prior $p(Z)$; i.e., the model is generative.

We argue that the deterministic models are still valuable:

the solutions are simple, cost-efficient, and easy to interpret. They can serve as nominal trajectories of human drivers in situations that can be approximated uni-modal. The nominal trajectories can be used in a trajectory planning algorithm which works with deterministic actors or disturbances.

For the scenarios with 5s prediction horizon, the time to compute a most-probable trajectory forecast is less than 10milliseconds on a single-core personal laptop with i7-6500U 2.50GHz CPU, and 8GB RAM without utilizing a parallelization. However, it takes several seconds (5-10s for 1,000 rollout trajectories) to construct the pdf for the probabilistic forecasts on the same machine due. One can significantly reduce the time via parallel computing (GPU).

V. Conclusion

Our work is the first attempt in the community to comprehensively understand and identify the impact of traffic signals on trajectory forecasting near TLs. In this regard, we first introduced the motivational example in Fig. 2 and defined 6 distinct scenarios of the problem. We proposed a novel idea to solve the scenarios that leverages the **future states** of traffic lights obtained via V2I communications. Specifically, we proposed deterministic and probabilistic human policy models to simulate state-dependent driver actions near TLs. In the ablation study, we show that the utilization of future phases and timings of TLs significantly improves the quality of trajectory forecasts for all scenarios described in Table. I.

As no dataset is publicly available that has the detailed TL data, interactions with surrounding vehicles, and the vision data, our experiments were based on a non-vision dataset with or without a front vehicle in the scene. Hence, a direct comparison against state-of-the-art forecasting models [1], [2], [3], [7], [8] is not made since the models were built on the vision data and/or with the access to all vehicles adjacent to the host vehicle (e.g., side or rear cars). Regardless, the proposed idea and the proposed framework can be leveraged by any work that concerns vehicle trajectory forecasting near TLs given the access to V2I communications. As the results suggest that the utilization of the future TL states could lead to the significant improvements in the prediction accuracy, we believe that it is worth building a large-scale dataset with both the TL and vision data to quantify the influence to the fullest extent.

Acknowledgment

References

[1] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker, "Desire: Distant future prediction in dynamic scenes with interacting agents," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 336–345.

[2] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2255–2264.

[3] J. Li, H. Ma, and M. Tomizuka, "Interaction-aware multi-agent tracking and probabilistic behavior prediction via adversarial learning," *arXiv preprint arXiv:1904.02390*, 2019.

[4] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 1468–1476.

[5] X. Li, X. Ying, and M. C. Chuah, "Grip: Graph-based interaction-aware trajectory prediction," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 3960–3966.

[6] R. Chandra, T. Guan, S. Panuganti, T. Mittal, U. Bhattacharya, A. Bera, and D. Manocha, "Forecasting trajectory and behavior of road-agents using spectral clustering in graph-lstms," *arXiv preprint arXiv:1912.01118*, 2019.

[7] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 961–971.

[8] R. Chandra, U. Bhattacharya, A. Bera, and D. Manocha, "Traphic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8483–8492.

[9] A. Bhattacharyya, M. Fritz, and B. Schiele, "Long-term on-board prediction of people in traffic scenes under uncertainty," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4194–4202.

[10] R. Chandra, U. Bhattacharya, C. Roncal, A. Bera, and D. Manocha, "Robusttp: End-to-end trajectory prediction for heterogeneous road-agents in dense traffic with noisy sensor inputs," in *ACM Computer Science in Cars Symposium*, 2019, pp. 1–9.

[11] C. R. Bennett and R. Dunn, "Driver deceleration behavior on a freeway in new zealand," *Transportation Research Record*, no. 1510, 1995.

[12] J. Wang, K. K. Dixon, H. Li, and J. Ogle, "Normal deceleration behavior of passenger vehicles at stop sign–controlled intersections evaluated with in-vehicle global positioning system data," *Transportation research record*, vol. 1937, no. 1, pp. 120–127, 2005.

[13] C. Sun, X. Shen, and S. Moura, "Robust optimal eco-driving control with uncertain traffic signal timing," in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 5548–5553.

[14] R. Akçelik and D. Biggs, "Acceleration profile models for vehicles in road traffic," *Transportation Science*, vol. 21, no. 1, pp. 36–54, 1987.

[15] G. Bham and R. Benekohal, "Development, evaluation, and comparison of acceleration models," in *81st Annual Meeting of the Transportation Research Board, Washington, DC*, vol. 6, 2002.

[16] ATS, "Acceleration/deceleration profiles at urban intersections," *Transit New Zealand Australasian Traffic Surveys*, 1990.

[17] P. P. Dey, S. Nandal, and R. Kalyan, "Queue discharge characteristics at signalised intersections under mixed traffic conditions," *European Transport*, vol. 55, no. 7, pp. 1–12, 2013.

[18] J. Park, D. Li, Y. L. Murphey, J. Kristinsson, R. McGee, M. Kuang, and T. Phillips, "Real time vehicle speed prediction using a neural network traffic model," in *The 2011 International Joint Conference on Neural Networks*. IEEE, 2011, pp. 2991–2996.

[19] B. Jiang and Y. Fei, "Vehicle speed prediction by two-level data driven models in vehicular networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 7, pp. 1793–1801, 2016.

[20] K. Fadhloun, H. Rakha, and P. Eng, "A vehicle dynamics model for estimating typical vehicle accelerations 2," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 35, no. 36, p. 37, 2015.

[21] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical review E*, vol. 62, no. 2, p. 1805, 2000.

[22] P. G. Gipps, "A behavioural car-following model for computer simulation," *Transportation Research Part B: Methodological*, vol. 15, no. 2, pp. 105–111, 1981.

[23] C. M. Bishop, "Mixture density networks," Citeseer, Tech. Rep., 1994.

[24] T. J. Gates, D. A. Noyce, L. Laracuente, and E. V. Nordheim, "Analysis of driver behavior in dilemma zones at signalized intersections," *Transportation Research Record*, vol. 2030, no. 1, pp. 29–39, 2007.

[25] G. Oh, D. J. Leblanc, and H. Peng, "Vehicle energy dataset (ved), a large-scale dataset for vehicle energy consumption research," *arXiv preprint arXiv:1905.02081*, 2019.

[26] D. Bezzina and J. R. Sayer, "Safety Pilot: Model Deployment Test Conductor Team Report," Tech. Rep. June, 2015. [Online]. Available: http://safetypilot.umtri.umich.edu/

[27] R. W. Schafer *et al.*, "What is a savitzky-golay filter," *IEEE Signal processing magazine*, vol. 28, no. 4, pp. 111–117, 2011.