

A Gray Box Model for Characterizing Driver Behavior

Soyeon Jung,¹ Ransalu Senanayake,² and Mykel J. Kochenderfer¹

¹ Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305

² Department of Computer Science, Stanford University, Stanford, CA 94305
{soyeonj, ransalu, mykel}@stanford.edu

Abstract

Understanding the behavior of human drivers is critical for the successful deployment of autonomous vehicles. Driver modeling is challenging due to the uncertainty inherent in human behavior and the complex interactions among drivers. Recent efforts have focused on black-box models that are highly expressive but lack interpretability of underlying dynamics. White-box models, on the other hand, are more interpretable as they are typically defined by relatively simple rules. They can prevent undesirable outcomes but cannot model the variability of human behavior. This paper presents a gray-box driver model that combines rule-based models with data-driven learning. The parameters of rule-based driver models are learned from real-world data using expectation-maximization. We perform experiments on interactive driving scenarios with lane changes and evaluate our model based on prediction accuracy, data efficiency, and safety. Results show that our model can accurately replicate human driver behavior with less data.

Introduction

Developing a realistic driver model that can capture human driving behavior is essential for the successful integration of autonomous vehicles into existing roadways. Autonomous vehicles, being safety-critical and complex systems, are required to comply with automotive safety standards. Traditionally, ISO 26262 has been the standard for functional safety of road vehicles (International Organization for Standardization 2011). It follows the principle that safety cannot be absolute, giving rise to the concept of tolerable risk. More recently, ISO/PAS 21448 was devised as the safety standard for driver assistance systems (International Organization for Standardization 2019). It accounts for any potential hazards despite the absence of system failures, and provides guidance to achieve the safety of the intended functionality (SOTIF). For validating safety-critical systems such as autonomous vehicles, simulations can be useful as they allow for testing of various scenarios without risking injury or death. However, the simulations need to employ realistic driver models in order to accurately evaluate the autonomous vehicles.

Various approaches have been explored to model human drivers. Recent studies have focused on black-box models including recurrent neural networks (Morton, Wheeler, and Kochenderfer 2016; Alahi et al. 2016; Zyner, Worrall, and Nebot 2019), generative adversarial networks (Gupta et al. 2018; Kosaraju et al. 2019), variational autoencoders (Ivanovic et al. 2020), and generative adversarial imitation learning (Kuefler et al. 2017; Bhattacharyya et al. 2019, 2020a). Although these expressive models can learn complex behaviors from data, they not only lack interpretability of the underlying dynamics but also can result in undesirable behaviors such as collision (Bhattacharyya et al. 2021).

On the other end of the spectrum, researchers have developed white-box models, or rule-based models, where the driver response is governed by a small set of predefined rules. White-box models, such as IDM (Treiber, Hennecke, and Helbing 2000) for car following and MOBIL (Kesting, Treiber, and Helbing 2007) for lane changing, are interpretable and thus able to prevent unacceptable behavior. However, they cannot model the inherent stochasticity of human behavior. To mitigate the limitations of both black-box and white-box models, Bhattacharyya et al. (2020b, 2021) developed a gray-box model where the parameters of an interpretable rule-based model are learned as probability distributions from data. For one-dimensional driving scenarios with no lane change, they estimate the parameters of the stochastic IDM (Treiber and Kesting 2017) using particle filtering. Because this model learns the underlying distributions by repeatedly sampling a set of particles, its computational complexity increases with the number of samples.

In this paper, we first extend the work of Bhattacharyya et al. (2021) to a two-dimensional driver model. The driver behavior is governed by the combination of a car-following model and a lane-changing model. Incorporating the lane-changing model adds more complexity and uncertainty to the interactions among multiple vehicles. We also propose a methodology to learn the parameters of the two-dimensional driver model from data using expectation-maximization (EM). One advantage of using EM is that it allows us to use expert knowledge to define the distributions over the parameters and measurements. Additionally, it is guaranteed that the likelihood monotonically increases with each iteration (McLachlan and Krishnan 2007). In this paper, we validate the hypothesis that our model can improve the accuracy and

data efficiency by performing experiments using a real world vehicle trajectory dataset.

Rule-based Driver Models

The driver behavior can be decomposed into two components: the longitudinal motion and the lane change behavior, which governs the lateral motion. This paper adopts a stochastic extension of Intelligent Driver Model (IDM) (Treiber and Kesting 2017) to model the longitudinal motion, and the Minimizing Overall Braking Induced by Lane change (MOBIL) model (Kesting, Treiber, and Helbing 2007) for the lane change behavior.

Stochastic IDM

The IDM (Treiber, Hennecke, and Helbing 2000) determines the longitudinal acceleration so the vehicle drives at the desired speed while maintaining safe separation from the vehicle in front. The input variables into the model are the absolute speed of ego vehicle $\dot{x}(t)$, the relative speed with respect to its preceding vehicle $\Delta\dot{x}(t)$, and the distance gap between two vehicles $d(t)$, at the current timestep t . The acceleration \ddot{x}_{IDM} is

$$\ddot{x}_{\text{IDM}} = a_{\text{max}} \left[1 - \left(\frac{\dot{x}(t)}{v_{\text{des}}} \right)^4 - \left(\frac{d_{\text{des}}}{d(t)} \right)^2 \right], \quad (1)$$

where d_{des} is the desired distance gap given by

$$d_{\text{des}} = d_{\text{min}} + \tau\dot{x}(t) - \frac{\dot{x}(t)\Delta\dot{x}(t)}{2\sqrt{a_{\text{max}}b}}. \quad (2)$$

The output variables d_{des} and \ddot{x}_{IDM} are governed by several parameters. The desired speed is v_{des} . The minimum allowable distance gap is d_{min} and the desired time gap to the preceding vehicle is τ . The limits on the acceleration and deceleration are a_{max} and b_{safe} , respectively.

In order to encode the inherent stochasticity in human driving behavior, we use the stochastic IDM (sIDM) (Treiber and Kesting 2017), which introduces an additional variance term σ_{IDM} . In this model, the acceleration output for each vehicle is sampled from the following Gaussian distribution with the mean \ddot{x}_{IDM} given in (1):

$$\ddot{x}_{\text{sIDM}} \sim \mathcal{N}(\ddot{x}_{\text{IDM}}, \sigma_{\text{IDM}}^2) \quad (3)$$

MOBIL

The MOBIL (Kesting, Treiber, and Helbing 2007) model determines whether to make a lane change in order to maximize the longitudinal acceleration of the ego vehicle and its neighbors. It initiates a lane change maneuver if the following conditions are met:

$$\tilde{\ddot{x}}_{\text{ego}} - \ddot{x}_{\text{ego}} + p(\tilde{\ddot{x}}_{\text{new}} - \ddot{x}_{\text{new}} + \tilde{\ddot{x}}_{\text{old}} - \ddot{x}_{\text{old}}) > \Delta a_{\text{th}} \quad (4)$$

$$-\tilde{\ddot{x}}_{\text{new}} \leq b_{\text{safe}} \quad (5)$$

In (4), the quantities with tildes are calculated assuming a lane change. Subscripts \cdot_{ego} , \cdot_{old} , and \cdot_{new} are associated with the ego vehicle, the follower before changing lane, and the follower after changing lane, respectively. The parameter $p \in [0, 1]$ is the politeness factor, which represents how

much the ego vehicle values the acceleration increase of its neighbors. The threshold of acceleration increase is Δa_{th} . The model decides to change lane only if a weighted sum of the acceleration increase of the ego vehicle and that of the neighbors exceeds this threshold. Equation (5) is a safety criterion to ensure that, if the lane change is made, the deceleration of the following car does not exceed the safe braking limit b_{safe} .

Vehicle Dynamics

Given the acceleration output from (3) and the lane-changing action from (4) and (5), the velocity and position of each vehicle can be updated according to the following dynamics. The longitudinal velocity and position for the next time step are

$$\dot{x}_{t+1} = \dot{x}_t + \ddot{x}_{\text{sIDM}}\Delta t \quad (6)$$

$$x_{t+1} = x_t + \dot{x}_t\Delta t + \frac{1}{2}\ddot{x}_{\text{sIDM}}\Delta t^2. \quad (7)$$

where \ddot{x}_{sIDM} is the sampled acceleration output. Unless a lane change is initiated, all vehicles maintain their lateral position along their path. Once a vehicle starts changing lane, the lateral movement is controlled by a PD controller (Kesting, Treiber, and Helbing 2007) until it reaches the center of the destination lane.

Driver Modeling using Expectation-Maximization

Throughout this section, $\mathbf{x} = \{(x^{(i)}, y^{(i)})\}_{i=1}^N$ denotes a set of known variables, that is, the longitudinal and lateral position measurements for each vehicle. We use \mathbf{z} to denote a vector of latent variables, in particular, the set of unknown IDM and MOBIL parameters for each vehicle. Our objective is to infer these latent parameters \mathbf{z} by observing \mathbf{x} . In this paper, we assume a discrete distribution over \mathbf{z} parameterized by θ .

Problem Formulation

Assuming \mathbf{z} is given, the probability of the i th vehicle changing lanes at each timestep is defined in a manner similar to MOBIL with

$$f_{\text{lane}}^{(i)}(\mathbf{x}, \mathbf{z}) = \begin{cases} \frac{1}{1+e^{-\lambda^{(i)}C}}, & \text{if (5) is met} \\ 0, & \text{otherwise} \end{cases}$$

where

$$C = \tilde{\ddot{x}}_{\text{ego}}^{(i)} - \ddot{x}_{\text{ego}}^{(i)} + p^{(i)} \left(\tilde{\ddot{x}}_{\text{new}}^{(i)} - \ddot{x}_{\text{new}}^{(i)} + \tilde{\ddot{x}}_{\text{old}}^{(i)} - \ddot{x}_{\text{old}}^{(i)} \right) - \Delta a_{\text{th}}^{(i)}. \quad (8)$$

Note that a sigmoid function is used to map the real-valued C to the range of 0 and 1. This replaces the hard constraint in (4) with a soft constraint, allowing us to incorporate stochasticity of the lane changing behavior into our model. We also introduce a new parameter λ which governs the degree of preference for switching lanes.

The likelihood of the next position $\mathbf{x}'^{(i)} = (x'^{(i)}, y'^{(i)})$ is obtained by the weighted sum of two cases, the lane-changing case and the car-following case, as follows:

$$p(\mathbf{x}'^{(i)} | \mathbf{x}, \mathbf{z}) = p_{\text{change}}(\mathbf{x}'^{(i)} | \mathbf{x}, \mathbf{z}) f_{\text{lane}}^{(i)} + p_{\text{follow}}(\mathbf{x}'^{(i)} | \mathbf{x}, \mathbf{z}) \left(1 - f_{\text{lane}}^{(i)}\right) \quad (9)$$

Parameter Estimation using EM Algorithm

Our goal is to find θ that maximizes the log-likelihood of the observations \mathbf{x} . For the i th vehicle, the log-likelihood can be written as

$$l(\theta^{(i)}) = \log \prod_{t=1}^{T_i} p(\mathbf{x}_t^{(i)}; \theta) \quad (10)$$

$$= \sum_{t=1}^{T_i} \log p(\mathbf{x}_t^{(i)}; \theta) \quad (11)$$

$$= \sum_{t=1}^{T_i} \log \sum_{\mathbf{z}_t^{(i)}} p(\mathbf{x}_t^{(i)}, \mathbf{z}_t^{(i)}; \theta) \quad (12)$$

$$= \sum_{t=0}^{T_i-1} \log \sum_{\mathbf{z}_t^{(i)}} p(\mathbf{x}_t'^{(i)} | \mathbf{z}_t^{(i)}, \mathbf{x}_t^{(i)}) p(\mathbf{z}_t^{(i)}; \theta). \quad (13)$$

This holds from the properties of logarithms (11), the law of total probability (12), and the conditional probability (13). However, due to the summation over the latent variable $\mathbf{z}_t^{(i)}$ inside the logarithm, $\theta^{(i)}$ cannot be solved analytically. Accordingly, the parameters $\theta^{(i)}$ are estimated based on EM algorithm (Dempster, Laird, and Rubin 1977). It constructs a tractable lower bound that contains a sum of logarithms as follows:

$$l(\theta^{(i)}) = \sum_{t=0}^{T_i-1} \log \sum_{\mathbf{z}_t^{(i)}} Q_t(\mathbf{z}_t^{(i)}) \frac{p(\mathbf{x}_t'^{(i)} | \mathbf{z}_t^{(i)}, \mathbf{x}_t^{(i)}) p(\mathbf{z}_t^{(i)}; \theta)}{Q_t(\mathbf{z}_t^{(i)})} \quad (14)$$

$$\geq \sum_{t=0}^{T_i-1} \sum_{\mathbf{z}_t^{(i)}} Q_t(\mathbf{z}_t^{(i)}) \log \frac{p(\mathbf{x}_t'^{(i)} | \mathbf{z}_t^{(i)}, \mathbf{x}_t^{(i)}) p(\mathbf{z}_t^{(i)}; \theta)}{Q_t(\mathbf{z}_t^{(i)})} \quad (15)$$

where the inequality in (15) holds from Jensen's inequality and log-concavity. This holds with equality if and only if $Q_t(\mathbf{z}_t^{(i)}) = p(\mathbf{z}_t^{(i)} | \mathbf{x}_t'^{(i)}, \mathbf{x}_t^{(i)}; \theta)$

The EM maximizes $l(\theta^{(i)})$ by iterating E-step and M-step until convergence. In the E-step, we compute $Q_i(\mathbf{z}_t^{(i)})$, the posterior probabilities of $\mathbf{z}_t^{(i)}$'s given \mathbf{x}_t and θ . For each timestep t ,

$$Q_t(\mathbf{z}_t^{(i)}) := p(\mathbf{z}_t^{(i)} | \mathbf{x}_t'^{(i)}, \mathbf{x}_t^{(i)}; \theta) = \frac{p(\mathbf{x}_t'^{(i)} | \mathbf{z}_t^{(i)}, \mathbf{x}_t^{(i)}) p(\mathbf{z}_t^{(i)}; \theta)}{\sum_{\mathbf{z}_t^{(i)}} p(\mathbf{x}_t'^{(i)} | \mathbf{z}_t^{(i)}, \mathbf{x}_t^{(i)}) p(\mathbf{z}_t^{(i)}; \theta)} \quad (16)$$

Then in the M-step, we find the maximum likelihood estimates for θ based on $Q_i(\mathbf{z}_t^{(i)})$ from the E-step:

$$\theta^* := \arg \max_{\theta'} \sum_{i=1}^n \sum_{t=0}^{T_i-1} \sum_{\mathbf{z}_t^{(i)}} Q_i(\mathbf{z}_t^{(i)}; \theta) \log \frac{p(\mathbf{x}_t'^{(i)} | \mathbf{z}_t^{(i)}, \mathbf{x}_t^{(i)}) p(\mathbf{z}_t^{(i)}; \theta')}{Q_i(\mathbf{z}_t^{(i)}; \theta)} \quad (17)$$

Experiments

In this section, we evaluate the performance of our trained model on a real-world dataset. Using our EM approach, we learn a distribution over four parameters: the desired longitudinal speed v_{des} , the stochasticity parameter in sIDM (σ_{IDM}), the politeness (p), and the lane-changing parameter (λ). We define each model parameter to follow a multinomial distribution, which allows us to work with discrete latent variables in the EM algorithm. For the other IDM/MOBIL parameters, we use the values provided in Table 1.

Once trained, we propagate simulated trajectories using the estimated parameters. Given a scenario with the initial position and velocity values of all vehicles, we iterate sampling the model parameters from the distributions and updating the position according to (7). Based on the simulated trajectories, we evaluate our model in terms of prediction accuracy, data efficiency, and safety.

Data Preprocessing

The proposed model is evaluated on a real-world dataset called the INTERnational, Adversarial and Cooperative motion (INTERACTION) dataset (Zhan et al. 2019). The dataset contains vehicle track data and roadway information data from highly interactive driving scenarios with cooperative and adversarial motions between drivers. In this work, experiments are conducted on a highway scenario with lane change and merging, shown in Fig. 1. We focus on the lane changing behavior and the lane following behavior.


Each entry of the vehicle track data consists of timestamp, track ID, position, velocity, orientation, vehicle length-

IDM parameter	Symbol	Value
Desired speed (m/s)	v_{des}	33.3
Desired time gap (s)	τ	1.5
Minimum acceptable gap (m)	d_{min}	2.0
Max acceleration (m/s^2)	a_{max}	1.4
Desired deceleration (m/s^2)	b	2.0
MOBIL parameter	Symbol	Value
Politeness	p	0.5
Safe braking (m/s^2)	b_{safe}	2.0
Acceleration threshold (m/s^2)	a_{th}	0.1

Table 1: IDM and MOBIL parameter values for the normal driver type (Kesting, Treiber, and Helbing 2009; Sunberg, Ho, and Kochenderfer 2017).



(a) A traffic scenario recorded from traffic cameras and drones (Zhan et al. 2019).



(b) A visualization of the processed vehicle tracks and roadway data.

Figure 1: A highway scenario with lane change and merging from the INTERACTION dataset (Zhan et al. 2019).

/width, and other information. We decompose the driver behavior into longitudinal and lateral motions with respect to the lanes by transforming the variables from Cartesian coordinates to Frenet coordinates. In Frenet coordinates, x represents the vehicle’s position (i.e. longitudinal displacement) along the reference path, and y represents side-to-side position (i.e. lateral displacement) relative to the reference path. We define the reference path as the centerline of the lane occupied by the vehicle.

Baselines

We compare the performance of our work against two baseline models. The first baseline is the IDM+MOBIL model using the default parameter values listed in Table 1. These values define the normal driver class as done in prior work (Kesting, Treiber, and Helbing 2009; Sunberg, Ho, and Kochenderfer 2017). In addition, the stochasticity parameter σ_{IDM} is set to zero, and the lane-changing actions are made deterministically based on (4) and (5). This model represents a purely rule-based, deterministic model.

The second model uses particle filtering (PF) to estimate the distribution over the latent variables. We adopt the method proposed by Bhattacharyya *et al.* (Bhattacharyya et al. 2021) to learn the MOBIL parameters as well as the IDM parameters. As in the EM approach, we use the default values in Table 1 for the parameters not being trained.

Prediction Accuracy

The prediction accuracy is measured using the discrepancy between the simulated trajectories and the ground truth tra-

jectories. We use the following two metrics introduced in (Gupta et al. 2018):

1. Average Displacement Error (ADE): The average ℓ_2 distance between the estimated points and the ground truth points over all time steps in the prediction horizon T . Then the value is averaged over n examples:

$$ADE = \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \sqrt{(\hat{x}_t^{(i)} - x_t^{(i)})^2 + (\hat{y}_t^{(i)} - y_t^{(i)})^2} \quad (18)$$

2. Final Displacement Error (FDE): The ℓ_2 distance between the estimated final position and the ground truth final position at the end of the prediction horizon T . Then the value is averaged over n examples:

$$FDE = \frac{1}{n} \sum_{i=1}^n \sqrt{(\hat{x}_T^{(i)} - x_T^{(i)})^2 + (\hat{y}_T^{(i)} - y_T^{(i)})^2} \quad (19)$$

Fig. 2 shows the ADE and FDE values for each model over different prediction horizons. We observe that both the PF and EM models are superior to the default model. For a detailed comparison, we also report in Table 2 the ADE and FDE of each model for 5-second duration and 10-second duration. According to these metrics, the PF-based model slightly outperforms the EM-based model, but their performance are nearly equivalent. A possible reason for the difference in their performance is that the EM model makes a stronger assumption about the data than the PF model.




Figure 2: Average displacement error (ADE) and final displacement error (FDE) over different time horizons.

	ADE-5	ADE-10	FDE-5	FDE-10
Default	4.222	14.048	10.036	43.666
PF	1.527	5.809	4.144	17.459
EM	1.618	5.925	4.373	17.490

Table 2: Average displacement error (ADE) and the final displacement error (FDE) in meters for 5-second duration and 10-second duration.

Data Efficiency

To evaluate the data efficiency, we train the models using subsets of different sizes from our dataset and compare their accuracy performance. Table 3 shows the ADE of the PF and EM models using small, medium, and large datasets. The small dataset contains only the first 10 data points of each vehicle track. The medium dataset contains the first 50 data points. The large dataset is the same as the original dataset where each track has about 200 data points on average. This experiment was not conducted on the default model because it does not involve any learning. We observe that both PF and EM models perform as poorly as the default model using the small dataset. With the medium dataset, we see that the performance of the EM model improved more than the PF model, indicating it is more data efficient.

	small (10 data points)	medium (50 data points)	large (original)
Default	–	–	4.222
PF	4.113	3.381	1.527
EM	4.292	2.705	1.618

Table 3: Average displacement error in meters for 5-second duration using different sized subsets of dataset.

Frequency	collisions	hard brakes
Default	0.0000	0
PF	0.0059	0
EM	0.0032	0

Table 4: Frequency of collisions and hard brakes in simulated trajectories using our EM model and two baselines.

Safety

To evaluate the safety, we inspect the frequency of undesirable behaviors in the simulated trajectories. These behaviors include collisions and hard brakes. When a vehicle is following its lane, collisions are counted when the headway distance is less than the vehicle length. During lane-changing period, collisions are counted when the ℓ_2 distance to another vehicle is less than 0.5 meters. Hard brakes are counted when a vehicle decelerates faster than the safe braking limit b_{safe} in (5).

Table 4 shows the frequency of collisions and hard brakes observed in the simulated trajectories for each model. The default model does not produce any collisions or hard brakes, as both IDM and MOBIL are designed to be collision-free. However, it suffers from poor prediction accuracy as seen in Fig. 2. The PF and EM based model also produce no hard braking. This is because in (8), we defined the probability to be positive only when (5) is satisfied. Collisions are observed in both PF and EM based model because they are probabilistic models. It turns out that EM based model achieves almost half the collision rate of the PF based model.

Conclusion

This paper presented a methodology for modeling human driver behavior that can efficiently learn a driver model without sacrificing safety. Our gray-box model can learn the variability from large amounts of data available, and at the same time, interpret the underlying dynamics of driving behavior. We performed experiments on a real-world driving scenario with lane changes and compared the performance our EM based model with two baselines, the default IDM+MOBIL model and the particle filtering based model. It was shown that our model can generate trajectories that represent the human driving driver in the real scenarios. The EM-based model was able to achieve nearly equal prediction accuracy to the PF based model with less data.

There are several potential directions for future work. While we assumed the latent variables of the EM algorithm

follow multinomial distributions, future work can use different distributions with a weaker assumption to better represent the actual distributions of the model parameters. In addition, we can evaluate the algorithms on other datasets such as Next-Generation Simulation (NGSIM) (Colyar and Halkias 2007). In addition, we can analyze the generalizability of the EM approach by applying it to other scenarios including merging and roundabout.

References

- Alahi, A.; Goel, K.; Ramanathan, V.; Robicquet, A.; Fei-Fei, L.; and Savarese, S. 2016. Social LSTM: Human trajectory prediction in crowded spaces. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 961–971.
- Bhattacharyya, R.; Jung, S.; Kruse, L. A.; Senanayake, R.; and Kochenderfer, M. J. 2021. A Hybrid Rule-Based and Data-Driven Approach to Driver Modeling Through Particle Filtering. *IEEE Transactions on Intelligent Transportation Systems*.
- Bhattacharyya, R.; Wulfe, B.; Phillips, D.; Kuefler, A.; Morton, J.; Senanayake, R.; and Kochenderfer, M. 2020a. Modeling human driving behavior through generative adversarial imitation learning. *arXiv preprint arXiv:2006.06412*.
- Bhattacharyya, R. P.; Phillips, D. J.; Liu, C.; Gupta, J.; Driggs-Campbell, K.; and Kochenderfer, M. J. 2019. Simulating emergent properties of human driving behavior using reward augmented multi-agent imitation learning. In *IEEE International Conference on Robotics and Automation (ICRA)*.
- Bhattacharyya, R. P.; Senanayake, R.; Brown, K.; and Kochenderfer, M. J. 2020b. Online parameter estimation for human driver behavior prediction. In *American Control Conference (ACC)*, 301–306.
- Colyar, J.; and Halkias, J. 2007. US highway 101 dataset. Technical Report FHWA-HRT-07-030, Federal Highway Administration (FHWA).
- Dempster, A. P.; Laird, N. M.; and Rubin, D. B. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1): 1–22.
- Gupta, A.; Johnson, J.; Fei-Fei, L.; Savarese, S.; and Alahi, A. 2018. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2255–2264.
- International Organization for Standardization. 2011. *ISO 26262: Road vehicles – Functional safety*. Geneva, Switzerland: International Organization for Standardization.
- International Organization for Standardization. 2019. *ISO/PAS 21448:2019: Road vehicles – Safety of the intended functionality*. Geneva, Switzerland: International Organization for Standardization.
- Ivanovic, B.; Leung, K.; Schmerling, E.; and Pavone, M. 2020. Multimodal deep generative models for trajectory prediction: A conditional variational autoencoder approach. *IEEE Robotics and Automation Letters*, 6(2): 295–302.
- Kesting, A.; Treiber, M.; and Helbing, D. 2007. General lane-changing model MOBIL for car-following models. *Transportation Research Record*, 1999(1): 86–94.
- Kesting, A.; Treiber, M.; and Helbing, D. 2009. Agents for traffic simulation. *Multi-agent Systems: Simulation and Applications*, 325–356.
- Kosaraju, V.; Sadeghian, A.; Martín-Martín, R.; Reid, I.; Rezatofghi, S. H.; and Savarese, S. 2019. Social-BiGAT: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. *arXiv preprint arXiv:1907.03395*.
- Kuefler, A.; Morton, J.; Wheeler, T.; and Kochenderfer, M. 2017. Imitating driver behavior with generative adversarial networks. In *IEEE Intelligent Vehicles Symposium (IV)*, 204–211.
- McLachlan, G. J.; and Krishnan, T. 2007. *The EM algorithm and extensions*, volume 382. John Wiley & Sons.
- Morton, J.; Wheeler, T. A.; and Kochenderfer, M. J. 2016. Analysis of recurrent neural networks for probabilistic modeling of driver behavior. *IEEE Transactions on Intelligent Transportation Systems*, 18(5): 1289–1298.
- Sunberg, Z. N.; Ho, C. J.; and Kochenderfer, M. J. 2017. The value of inferring the internal state of traffic participants for autonomous freeway driving. In *American Control Conference (ACC)*, 3004–3010.
- Treiber, M.; Hennecke, A.; and Helbing, D. 2000. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2): 1805.
- Treiber, M.; and Kesting, A. 2017. The intelligent driver model with stochasticity—new insights into traffic flow oscillations. *Transportation Research Procedia*, 23: 174–187.
- Zhan, W.; Sun, L.; Wang, D.; Shi, H.; Clausse, A.; Naumann, M.; Kümmerle, J.; Königshof, H.; Stiller, C.; de La Fortelle, A.; and Tomizuka, M. 2019. INTERACTION Dataset: An INTERnational, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps. *arXiv:1910.03088 [cs, eess]*.
- Zyner, A.; Worrall, S.; and Nebot, E. 2019. Naturalistic driver intention and path prediction using recurrent neural networks. *IEEE Transactions on Intelligent Transportation Systems*, 21(4): 1584–1594.