

# Hierarchical Learned Risk-Aware Planning Framework for Human Driving Modeling

Nathan Ludlow<sup>1</sup> Yiwei Lyu<sup>2</sup> and John Dolan<sup>2</sup>

**Abstract**—This paper presents a novel approach to modeling human driving behavior, designed for use in evaluating autonomous vehicle control systems in a simulation environments. Our methodology leverages a hierarchical forward-looking, risk-aware estimation framework with learned parameters to generate human-like driving trajectories, accommodating multiple driver levels determined by model parameters. This approach is grounded in multimodal trajectory prediction, using a deep neural network with LSTM-based social pooling to predict the trajectories of surrounding vehicles. These trajectories are used to compute forward-looking risk assessments along the ego vehicle’s path, guiding its navigation. Our method aims to replicate human driving behaviors by learning parameters that emulate human decision-making during driving. We ensure that our model exhibits robust generalization capabilities by conducting simulations, employing real-world driving data to validate the accuracy of our approach in modeling human behavior. The results reveal that our model effectively captures human behavior, showcasing its versatility in modeling human drivers in diverse highway scenarios.

**Index Terms**—Human modeling; Risk-aware; Social Pooling; Trajectory Prediction; Autonomous Vehicles; Simulation; Robotics

## I. INTRODUCTION

As the presence of autonomous vehicles on our roads becomes more common, rigorous testing of safety algorithms is essential for the protection of all road users. Real-world testing poses numerous challenges, including cost, time, and safety concerns, particularly in scenarios with potential near-collision events. While real-world validation of systems remains essential, simulations offer a compelling alternative, providing safety, efficiency, cost-effectiveness, and the ability to test and optimize systems under various challenging scenarios without physical risk.

However, a significant challenge in simulation testing stems from the unpredictable nature of human drivers, who often cause near-collision situations involving autonomous vehicles. To ensure accurate and effective simulation testing of autonomous vehicle systems, it’s crucial to integrate a precise model of human driving behavior. This ensures that algorithms can adeptly handle real-world unpredictability, ultimately enhancing road safety.

Despite numerous proposed approaches for modeling human behavior, no widely adopted comprehensive model of human driving behavior exists, largely due to limitations in existing models. Human drivers typically plan their actions

many seconds in advance, considering numerous driving factors and potential actions of surrounding vehicles, such as lane selection, overtaking decisions, and yielding to oncoming traffic [1]. To accurately replicate human-like driving behavior for advanced maneuvers, we propose the need to plan the ego vehicle’s trajectory several seconds into the future. This planning must account for the multimodal nature of the behavior of surrounding vehicles while optimizing for a safe trajectory, considering all potential surrounding vehicle actions. Many existing human-driving models have limited planning horizons and single-mode surrounding vehicle predictions, which hinder their ability to execute higher-level human behaviors and navigate safely around multiple vehicles. Leveraging an intuitive understanding of how human drivers plan their vehicle trajectories, our proposed model offers improved capabilities in modeling human driving behavior.

To achieve this, we propose a novel hierarchical learned risk-aware planning framework. Our framework uses a recurrent neural network-based trajectory prediction with social pooling method to predict the actions of all surrounding vehicles. Then to capture the interactive risk the ego vehicle faces we propose a matrix to represent the aggregated risk at each point along the highway from a multimodal distribution of possible trajectories from all surrounding vehicles. We use Next Generation Simulation (NGSIM) [2] data to learn parameters that determine human driving and feed those into the risk-aware planning cost function to more accurately represent human driving based on these learned parameters. Since our proposed model is able to plan a long-term trajectory in real-time, this framework is effective at capturing high-level human behaviors, while being sufficiently computationally to be used in simulations with a high number of agents. Additionally, the model includes a parameter for maximum allowable risk, allowing the model to be adjusted to represent different styles of driving, increasing the accuracy of the model in various situations.

Our **main contributions** are: (1) A novel forward-looking risk metric based on the hierarchical combination of social pooling and history-based trajectory prediction and risk evaluation along predicted trajectories that evaluates risk based not only on current vehicle states, but also predicted future vehicle states that are estimated based on vehicle history and vehicle interaction effects; (2) An accurate human model adjustable with parameters found from data-driven approaches that can represent multiple styles of driving behaviors while being computationally efficient, allowing it to be used in real-time risk assessment and planning of autonomous vehicles. (3) An assessment of the hierarchical learned risk-aware

<sup>1</sup>Robotics and Dynamics Laboratory, Brigham Young University, Provo, Utah 84602, USA ndl22@byu.edu

<sup>2</sup>CMU Center for Autonomous Vehicle Research, Carnegie Mellon University, Pittsburgh 15289, USA [yiweilyu, jdolan]@andrew.cmu.edu

planning framework’s performance using real-world data, coupled with simulation-based examinations of its capacity to replicate human driving behavior, accompanied by an analysis of human driver conduct within specific contextual scenarios.

## II. RELATED WORK

Rule-based approaches attempt to model human driving behavior by intuiting how a human driver will maneuver their vehicle and defining a comprehensive set of rules to represent that behavior. These approaches often generalize well and can easily generate new behaviors not previously observed in datasets. These models also have the benefit of generating easily interpretable trajectories. There are many rule-based approaches for modeling human behaviors such as the Intelligent Driver Model (IDM) [3], Game-theoretic planning [4], [5], and risk-aware planning [6], [7]. While the IDM works well in its context, it can only replicate lane following without lane changing, making it unable to replicate complex human behavior. Other approaches can be used to build comprehensive driving models, but fail to driving behavior that is human-like [8].

Learning-based approaches to modeling human driving behavior generally use Inverse Reinforcement Learning [9]–[11] or Imitation Learning [11]–[13] to learn cost functions for human drivers or to replicate expert driving maneuvers, respectively. Both of these methods perform reasonably at replicating human driving behavior when tested in situations already observed in data; however, these methods fail to generalize well to situations and contexts not seen in training data [14]. Additionally, learning-based methods require large amounts of data for training, which is not always available [15]. Another problem of these learning-based models is that due to the complexity of driving they fail to learn high-level behaviors of vehicles, and only effectively to capture the specific maneuvers of individual vehicles [14].

Human driving behavior depends on a number of factors; however, it has been shown that human drivers operate their vehicles in large part based on perceived risk [7], [16]–[18]. These studies however, do not attempt to use risk modeling to build a comprehensive model of human driving and instead only show that certain high level behaviors can be predicted from the risk the ego vehicle experiences. In our work, we assume that drivers operate their vehicle primarily based on risk [19]–[21], and perceived optimality to the driver of a path, and we propose that we can develop an effective driver model with this method. The NGSIM dataset provides valuable information on human driving behavior with proper filtering of data [22]–[24]. By using these data, we observe human driver tendencies and behaviors to learn and validate models on real-world data. Our novel approach to modeling human driving behavior has distinct advantages over existing methods in multiple key areas. By predicting the trajectories of surrounding vehicles multiple seconds into the future, we enable the risk-aware planner to plan for many maneuvers that may be less optimal in the short term but yield long-term rewards, similar to how human drivers often operate

their vehicles. Additionally, by learning parameters for the risk-aware planning we can gain many of the benefits of both rule-based and learned models by having a well defined set of driving rules that are generalized to perform well in all road contexts and situations even when the model has not been trained on the specific scenario, while still accurately representing human driving behavior. Lastly, by having adjustable parameters for maximum allowable risk our framework can represent different styles of driving based on the aggressiveness specified for the driver model.

## III. METHOD

### A. Trajectory Prediction

Convolutional Social Pooling for trajectory prediction is the means of estimating a set of probable trajectories for a vehicle [25]. The trajectory prediction is performed by a set of Long Short-Term Memory (LSTM) networks that are used to encode the histories of surrounding vehicles into a social occupancy grid. These histories are then passed through multiple convolutional and pooling layers before being decoded by another LSTM network into predicted trajectories. The predictor outputs six individual trajectories, representing three lateral maneuvers: left lane change, keep lane, and right lane change; and two longitudinal maneuvers: maintaining velocity or decelerating. Each of these trajectories is assigned an associated probability based on its likelihood of occurrence. The Convolutional Social Pooling approach from [25] is employed for its ability to model human driving behavior, its computational efficiency for real-time predictions, and its multimodal output. While Convolutional Social Pooling is well-suited for this purpose, other sufficiently accurate multimodal trajectory prediction methods could be adapted to fit this framework.

### B. Risk Measurement

Our approach considers all other agents to be dynamic obstacles with respect to the ego agent. We consider a bounded environment  $Q \subset \mathbb{R}^2$  with points  $q \in Q$ . The positions of  $N_{agents}$  agents in the environment are denoted as  $p_j$  with  $j = \{1, \dots, N_{agents}\}$ , and velocities of the agents are denoted with  $u_j$ . The position of the ego agent is denoted as  $p_e$  and the velocity as  $u_e$ . Each agent is modeled as a kinematic bicycle model with acceleration  $a_i$  and steering  $\theta_i$  the control input for agent  $j$ . To account for the uncertain multimodal nature of human driving we use the trajectory prediction module introduced above [25] to generate a series of potential trajectories for each actor in the scene at future time steps  $t_i$  spaced evenly  $t_s$  apart where  $i = \{1, \dots, N_{time}\}$  and  $N_{time}$  is the number of predicted vehicle positions. The prediction module can account for interaction effects of other vehicles as well as vehicle histories to provide a more accurate model of potential actions taken by surrounding vehicles. Additionally the trajectory prediction module accounts for multimodality in the potential human actions by generating  $N_{traj}$  possible trajectories with the associated probability  $P_{ik}$  with  $k = \{1, \dots, N_{traj}\}$ , indicating the normalized likelihood the vehicle will follow the associated

trajectory, bounded within the range of  $[0, 1]$ . Each trajectory is composed of a series of points  $p_{ijk}$  spaced  $\Delta t$  apart in time, indicating the expected position of the vehicle. This multimodal risk prediction allows us to plan the ego vehicle trajectory based on multiple probable futures for each vehicle similar to how human drivers operate their vehicles.

$$L_{\bar{P},i} = \{q \mid \mathcal{H}_{c,i}(q, p, \dot{p}) \leq \mathcal{H}_P\} \quad (1)$$

For each future time step  $t_i$  of the predicted future trajectories of each trajectory we represent the risk at each point  $q$  by defining a risk cost as  $\mathcal{H}_{c,i}$ . As in [7], the risk cost represents the risk posed to any point along the road relative to the positions and velocities of all surrounding agents. We define a planning threshold as  $\mathcal{H}_P$  as the maximum risk the ego vehicle navigate through. This leads to the definition of a risk level set  $L_{\bar{P},i}$  represented in Equation 1 at a future time step  $t_i$  which is the set of all points less than the chosen planning threshold  $\mathcal{H}_P$  that the ego vehicle uses when planning.




Fig. 1. A visualization of the normalized risk computed using the higher-order rectangular Gaussian function for a single vehicle at one time step. Lower risk areas are visualized in blue, and higher risk areas in yellow.

$$\mathcal{H}_c(q, p_i, \dot{p}_i) = \sum_{i=1}^n \frac{\exp\left(-\left(\frac{(q_x - p_{x,i})^2}{\sigma_{x,i}^2}\right)^\beta - \left(\frac{(q_y - p_{y,i})^2}{\sigma_{y,i}^2}\right)^\beta\right)}{1 + \exp(\alpha \dot{p}_i^T (q - p_i))} \quad (2)$$

To compute  $\mathcal{H}_{c,i}$  we use rectangular higher-order Gaussians from the work in [7] represented in Equation 2 with  $\sigma_i = \frac{w_i}{2} + |\dot{p}_i|$  for agent  $i$  with  $w_i$  representing the size of the agent modeled as a rectangle. We choose to use the rectangular Gaussian as opposed to the elliptical variant, as it weights the risk equally along all points of the body of the agent. This approach allows our method to accurately compute the risk surrounding vehicles of any length or width along the road. Additionally the higher-order Gaussian method of computing risk takes into account both the position and velocity of each agent, meaning that a vehicle with a higher velocity will cause higher risk at a point than an equivalent vehicle with a lower velocity. This intuitively means that as other agent vehicles travel faster the ego vehicle will plan to stay further from them, similar to how human drivers operate their vehicles.

We modify the rectangular cost function to include the probability  $P_{jk}$  of the trajectory, since trajectories that are less probable to be taken by the surrounding vehicles cause less risk. As used in [7], we use  $\alpha = 0.8$  and  $\beta = 1.5$ , since this provided the best fit to the data based on the behavior from cars in the NGSIM data. By modifying the cost function to take into account the probability of the trajectory being computed we allow the ego vehicle to prioritize avoiding paths that would collide with highly probable trajectories of surrounding vehicles, while also planning to avoid paths that could collide with low likelihood trajectories of surrounding vehicles when there is sufficient space on the road available.

$$\mathcal{H}_{c,i}(q, p_i, \dot{p}_i) = \frac{\sum_{j=1}^{N_{agents}} \sum_{k=1}^{N_{traj}} \cdot P_{ij} \cdot \exp\left(-\left(\frac{(q_x - p_{x,ijk})^2}{\sigma_{x,ijk}^2}\right)^\beta - \left(\frac{(q_y - p_{y,ijk})^2}{\sigma_{y,ijk}^2}\right)^\beta\right)}{1 + \exp(\alpha \dot{p}_{ijk}^T (q - p_{ijk}))} \quad (3)$$

Equation 3 gives us the risk cost for all points  $q$  in the planning region at a given future time step  $t_i$ . We define  $\sigma_{ijk} = \frac{w_{ijk}}{2} + |\dot{p}_{ijk}|$  for the trajectory  $k$  from vehicle  $j$  at time step  $i$  accounting for the probability of each trajectory and its time step of occurrence. From the risk of all points at the given time we use Equation 1 to find the risk level set of points the ego vehicle can use to plan for that time step. With this approach we seek to find the risk the ego vehicle will experience at any point along the road over the next few seconds. With this approach we can plan paths that are multiple seconds long before determining which path to take. This is an important aspect of the framework, since human drivers operate their vehicles by thinking about the benefits or dangers or maneuvers over the timeframe of a few seconds. By planning for trajectories that span multiple seconds we can replicate this aspect of human driving and more select paths that imitate human driving behavior.

### C. Path Planning and Navigation

After computing the risk at each future time step  $t_i$  that is predicted for each of the other agents we generate a grid of points or nodes spaced  $d_p$  along the length of the road  $m = \{1, \dots, N_{horiz}\}$  and for each lane on the road where  $l = \{1, \dots, N_{lanes}\}$  with the position of each node represented as  $g_{ml}$ . We connect the nodes of this graph with edges that represent possible actions the car can take from any node. These actions are: continue in the current lane, change to the left lane, or change to the right lane. We assign weights  $w_{mlc}$  according to Equation 4 to each edge with  $c = 1, 2, 3$  for each of the possible lane change actions, and at each node we assign the corresponding level set  $L_{\bar{P},i}$  to each point in the planning grid  $L_{ml}$  for if the risk at that point is lower than the risk threshold  $\mathcal{H}_P$ .

$$w_{mlc} = \mathcal{H}_{c,i}(g_{ml}, p_i, \dot{p}_i) \quad (4)$$

Using the planning grid, we employ Dijkstra’s Algorithm to determine the lowest-overall-risk path for the ego vehicle similar to the approach from [7]. Dijkstra’s Algorithm is selected to guarantee that a globally optimal path is found for the ego vehicle, and works efficiently for this algorithm due to the relatively small number of nodes required for effective path for planning. At each point, we verify if  $L_{ml}$  falls within the planning level set  $L_{\hat{P},i}$  and incorporate it as a node in the planning graph only when within this level set. If a complete path cannot be established from start to finish, we select the longest path with the lowest cost. This approach minimizes overall risk while avoiding exceeding a maximum safety threshold, thus emulating various human behaviors that prioritize overall safety relative to other vehicles on the road. Thresholding max risk with  $L_{P,i}$  prevents the vehicle from taking single overly risky maneuvers that humans would avoid, even if they result in lower risk along the longer path, while optimizing for the lowest-cost total path ensures the vehicle’s safety throughout.




Fig. 2. The planning grid with nodes and edges for the ego vehicle (pictured in blue) that is used to find a safe path through the surrounding agents.

#### Algorithm 1 Generate Multimodal Risk-Aware Planning Grid

```

1: for  $\forall j \in \{1, \dots, N_{agents}\}$  do
2:   Predict possible agent trajectories [25]
3: end for
4: Create graph in planning horizon length
5: for  $\forall i \in \{1, \dots, N_{time}\}$  do
6:   Compute risk  $\mathcal{H}_{c,i}$  (Equation 3)
7: end for
8: for  $\forall m \in \{1, \dots, N_{horiz}\}$  do
9:   for  $\forall l \in \{1, \dots, N_{lanes}\}$  do
10:    Find time step  $i$  for ego vehicle to reach edge  $g_{ml}$ 
11:    Update  $w_{ml}$  for each edge (Equation 4)
12:    Set  $L_{ml}$  to the risk level set (Equation 1)
13:   end for
14: end for

```

#### D. Insights on Learned Parameters

To more accurately represent human driving behavior as opposed to simply generating safe driving behavior we take the framework previously outlined and learn additional costs that are used when planning trajectories. Not all lanes along a highway travel at the same speed: human drivers prefer to drive in different lanes based on their desired speed of travel. To reflect this behavior in our model, we learn the driver’s lane preference based on the current speed of the driver relative to the speed of other vehicles in their lane.

To learn parameters as additional costs for the risk-aware planning, we divide the NGSIM data from I-80 and US-101 in half, reserving half of the data to validate our model. We also filter all data to find the smoothed positions and velocities of all the vehicles to correct errors in the NGSIM

data. Finally, we re-extract the current lanes of each vehicle from its position to accurately identify each vehicle’s lane at each time step in the NGSIM data. The data used for training the model included approximately 4,500 cars with over 80 hours of combined driving time for learning parameters to model the driving on human behavior.

TABLE I  
LEARNED PARAMETERS FOR MODEL

	Lane 1 (Slowest)	Lane 2	Lane 3	Lane 4	Lane 5 (Fastest)
$1\sigma$ slower	0.4224	0.5030	0.5169	0.4994	0.4460
Mean	0.4017	0.4868	0.4862	0.4677	0.4386
$1\sigma$ faster	0.4123	0.4962	0.5070	0.4891	0.4421




Fig. 3. A visualization of the risk in each lane relative to the velocity of the vehicle. This is a visualization of data seen in Table I.

For each time step of each vehicle in the data, we compute the risk field at the current time step  $t_i$  using the rectangular cost function in Equation 2, and record the risk the vehicle is experiencing, along with its velocity, current lane, and the mean and standard deviation of velocity of all vehicles in the current vehicle’s lane. We process these data to determine the average risk a vehicle experiences when traveling within the range of  $1.5$  standard deviations from the mean velocity of the lane  $[-1.5\sigma, 1.5\sigma]$ . We selected  $1.5\sigma$  since about 90% of all vehicle velocities fall within this range, leading to a good fit for the data without giving too much weight to outliers. The data can be found in Table I (see Figure 3 for a visualization of these data).

To incorporate the learned parameters into our model we modify the weights of the edges of the planning grid with Equation 6 to reflect the difference in risk that the vehicle would experience by changing lane. We fit a function to the acquired data  $\mathcal{H}_l(l, u_{avg}, u_\sigma, u_{ijk})$  represented by Equation 5 with an input of the current current lane, average lane velocity, standard deviation of lane velocity and the velocity of a vehicle and output of the average risk in the lane bounded in the range  $[0, 1]$ . Then we compute the current value of  $\Delta\mathcal{H}_l$  using Equation 5.

$$\Delta\mathcal{H}_l = \mathcal{H}_l(l2, u_{l2,avg}, u_{l2,\sigma}, u_{i+1,jk}) - \mathcal{H}_l(l1, u_{l1,avg}, u_{l1,\sigma}, u_{ijk}) \quad (5)$$

$$w_{mlc} = \mathcal{H}_{c,i}(g_{ml}, p_{ijk}, \dot{p}_{ijk}) + \Delta\mathcal{H}_l \quad (6)$$

Using Equation 6 as the new cost function for the edges of the planning grid in Algorithm 1 generates two new behaviors in the ego vehicle that replicate human driving behavior in highway driving. First, it rewards the ego vehicle for traveling at the speed of the other traffic in the lane. As the ego vehicle speeds up or slows down to reach the same velocity as the other vehicles in the lane the  $\Delta\mathcal{H}_l$  factor decreases, which lowers the total cost of the path. This behavior of speeding up or slowing down in a lane to match the average speed of other cars in the lane more closely matches human behavior when driving within one lane on a highway as opposed to constantly speeding up and slowing down to maintain a safe distance other vehicles.

Second, the  $\Delta\mathcal{H}_l$  factor encourages the ego vehicle to enter lanes that align with the ego vehicle’s desired speed. We input the ego vehicle’s desired velocity as a planner parameter, which guides it to lanes with velocities matching its preference. When the ego vehicle aims to surpass its current lane’s speed, it shifts to a faster lane, lowering its  $\Delta\mathcal{H}_l$  term. Conversely, it changes to a slower lane when desiring a lower speed than the current lane’s average velocity. These emergent behaviors, arising from this single cost term addition, highlight the approach’s advantages. Many such terms can be learned from data to improve the modeling of human behaviors in varying driving situations. In the next section, we validate our model’s lane change prediction ability using vehicle data and demonstrate its effective use in replicating human driver lane changes. The process for learning parameters can be easily repeated for new parameters, continually generating accurate human behaviors for various driving scenarios.

#### IV. SIMULATION AND DISCUSSION

##### A. Model validation

In this section, we validate our model using real-world data and show specific examples of our framework generating human-like behavior not found in other simulation methods for human drivers. We validate the model using the approximately 1,000 lane change maneuvers in the half of the NGSIM data that we did not use to learn the parameters. Further, we subdivide the data by examining the lateral velocities of each driver during their maneuvers and categorizing the drivers as conservative, normal, or aggressive. We define drivers with the lower quartile of lateral and longitudinal accelerations during lane changes as conservative, drivers in the upper quartile as aggressive, and the rest as normal. We define three different parameters for aggressiveness that represent the varying driving styles, and validate each lane change maneuver with the corresponding model.

To perform the validation, we replace one of the recorded vehicles from NGSIM data in the simulation with our model 2 seconds before the lane change occurs in the data, and observe the planned trajectory for the next 4 seconds to determine if the model changes lanes into the same lane as the actual data. We implement the previously described method of planning with a model predictive controller and record the final and average deviation of our model from

the ground truth. We divide the error into lateral, and longitudinal error to demonstrate the models performance in both directions. We present our results in Table II.

TABLE II  
DRIVER MODEL VALIDATION BY DRIVER TYPE

	Aggressive	Normal	Conservative	Total
<b>Accuracy</b>	77.41%	81.27%	83.73%	80.92%
<b>RMSE<sub>lat,avg</sub></b>	0.348 m	0.312 m	0.285 m	0.315 m
<b>RMSE<sub>lat,final</sub></b>	0.518 m	0.436 m	0.384 m	0.446 m
<b>RMSE<sub>lon,avg</sub></b>	3.51 m	3.12 m	2.98 m	3.20 m
<b>RMSE<sub>lon,final</sub></b>	5.32 m	5.11 m	5.01 m	5.15 m

Table II demonstrates our model’s capability to accurately predict lane changes. Notably, prediction accuracy decreases slightly as driver aggressiveness increases. This is likely due to the more variable nature of aggressive driving behaviors. Aggressive drivers often prioritize speed as-opposed-to safety and adherence to traffic rules, expanding the set of potential actions beyond those of conservative drivers. This variability makes it challenging to pinpoint the exact actions of aggressive drivers, as reflected in our results. Importantly, the slight decrease in accuracy for aggressive drivers does not hinder our model’s capability to generate plausible aggressive driving maneuvers. The framework may select a slightly different, but realistic aggressive maneuver within the broader range of possibilities.

TABLE III  
COMPARISON OF EXISTING MODELS FOR PREDICTING LANE CHANGE MANEUVERS WITH 3 SECONDS OF PREDICTION HORIZON

Model	Dataset	Accuracy
Proposed	NGSIM	81%
Support-vector Machine [26]	NGSIM	47%
Multi-Layer Perceptrons [26]	NGSIM	58%
Random Forest [26]	NGSIM	63%
LSTM Network [27]	NGSIM	86%
Convolutional Neural Network [26]	NGSIM	91%

We provide a comparison of our model’s ability to predict lane changes by human drivers in Table III. While the Convolutional Neural Network and LSTM Network yield slightly better results in lane prediction than our proposed model, we assert that our model remains valid as a method for modeling human driving behavior as our model prioritizes generating new human-like trajectories over precise lane change predictions. Furthermore, we comprehensively model the entire highway driving scenario, providing insights into both the likelihood and location of lane changes, in contrast to the binary classification of lane change occurrence offered by many other methods. We additionally emphasize the uniqueness of our approach in modeling human driving behavior based on adjustable hyperparameters for aggressiveness. These parameters enable our model to generate multiple

distinct human-like trajectories for the same initial highway scenario leading to an improved capability of our model to be used for testing autonomous vehicle frameworks.

We acknowledge that training exclusively on NGSIM data restricts the model’s ability to reproduce driving maneuvers absent from the dataset. However, NGSIM’s diverse range of drivers suggests that, with sufficient data on other driving maneuvers, this approach can learn additional behaviors.

### B. Example Cases

To demonstrate the ability of this model to generate new human behaviors not seen in data we present the following scenarios which showcase our framework’s ability to generate diverse behaviors. All scenarios shown below are sequential snapshots of egocentric views of a simulated highway environment with the ego vehicle is depicted in dark blue and the planned trajectory shown by the pink line.

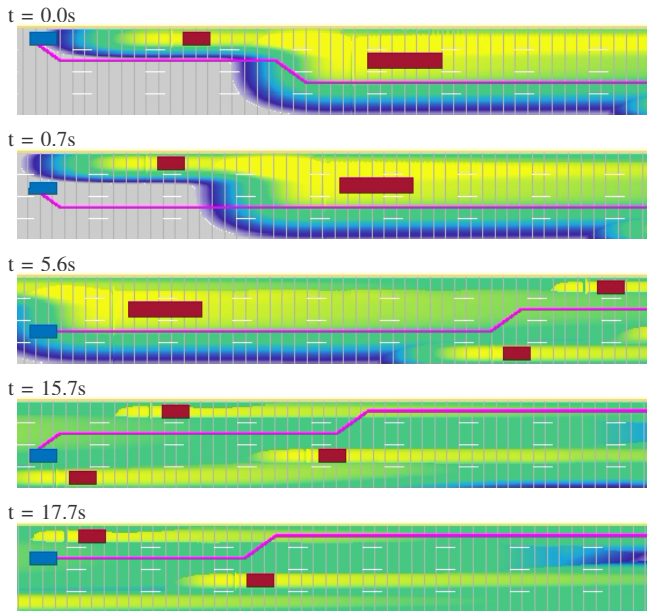


Fig. 4. A demonstration of the model’s ability to weave through traffic. The ego vehicle starts behind all vehicles on the road, but with a higher velocity and risk threshold  $\mathcal{H}_P$  than all surrounding vehicles. The risk-aware planner navigates a safe human-like route to weave through the traffic. [Video](#)

1) *Weaving through traffic:* In Figure 4, the ego vehicle, with a higher desired velocity and risk threshold  $\mathcal{H}_P$ , displays more aggressive behavior compared to its surroundings, weaving through traffic to emulate aggressive driving. Current human behavior models struggle in weaving through dense traffic due to two main limitations. Firstly, models with short planning horizons lack sufficient time to plan a path through dense traffic, causing the ego vehicle to follow dense traffic groups rather than weaving through them. Secondly, when existing models attempt weaving, they frequently collide with surrounding vehicles due to a lack of consideration for the multimodal nature of surrounding vehicle trajectories which leads to high levels of uncertainty in the actions of other vehicles. Our approach, with its long planning horizon and multimodal trajectory prediction, safely

navigates traffic, accurately representing aggressive driving behavior.

2) *Moving out of way of high-speed vehicle:* In Figure 5, the ego vehicle, with a lower desired velocity and risk threshold  $\mathcal{H}_P$ , exhibits a more cautious approach than surrounding vehicles as it encounters a fast-approaching car from behind. Our Risk-Aware planning framework guides the ego vehicle to change lanes, allowing the faster vehicle to pass. This behavior accurately models a conservative driver prioritizing safe driving by making way for a faster driver. Most current driver models lack this behavior due to inadequate forward-looking risk assessments for surrounding vehicles. Consequently, they fail to predict potential risks from other vehicles early enough, resulting in a keep-lane trajectory that causes the following vehicle to decelerate, even when a conservative human driver would change lanes. With our forward-looking risk assessment, the ego vehicle predicts the risk posed by the approaching vehicle well in advance, enabling it to yield the lane to the faster vehicle, mirroring the behavior of a cautious human driver.

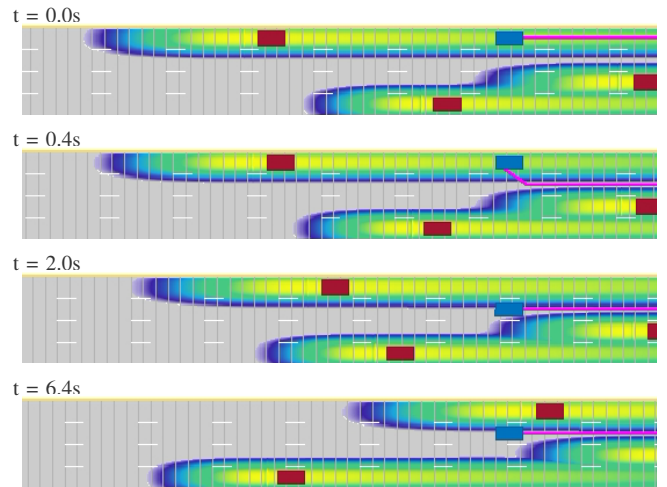


Fig. 5. A demonstration of the model’s ability to move out of the way of an approaching high-speed vehicle to maintain safety. The ego vehicle can be seen starting in front of a high-speed vehicle and traveling at a slower velocity and risk threshold  $\mathcal{H}_P$  than the approaching vehicle. The planner navigates a safe human-like route to move out of the way. [Video](#)

## V. CONCLUSION

In this paper, we propose the hierarchical learned risk-aware planning framework for modeling human drivers and validate and demonstrate its effectiveness in replicating human behavior. We additionally showcase the model’s ability to generate human-like driving patterns in diverse, unseen situations. Future work will involve enhancing the model to encompass different driving characteristics [28], [29], such as distracted driving, broadening the range of drivers it can represent. We also plan to use methods like probing [30] to actively learn additional parameters to improve the models ability to mimic human behavior across a wider array of situations and conduct further validation in diverse contexts to confirm its generalizability to other driving scenarios.

## REFERENCES

- [1] Y. Lyu, C. Dong, and J. M. Dolan, “Fg-gmm-based interactive behavior estimation for autonomous driving vehicles in ramp merging control,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020, pp. 1250–1255.
- [2] U.S. Department Of Transportation Federal Highway Administration, *Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data*, en, 2017.
- [3] M. Treiber, A. Hennecke, and D. Helbing, “Congested traffic states in empirical observations and microscopic simulations,” en, *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, Aug. 2000.
- [4] J. F. Fisac, E. Bronstein, E. Stefansson, D. Sadigh, S. S. Sastry, and A. D. Dragan, *Hierarchical Game-Theoretic Planning for Autonomous Vehicles*, arXiv:1810.05766 [cs, math], Oct. 2018.
- [5] H. Yu, H. E. Tseng, and R. Langari, “A human-like game theory-based controller for automatic lane changing,” en, *Transportation Research Part C: Emerging Technologies*, vol. 88, pp. 140–158, Mar. 2018.
- [6] A. Pierson, W. Schwarting, S. Karaman, and D. Rus, “Navigating Congested Environments with Risk Level Sets,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, Brisbane, QLD: IEEE, May 2018, pp. 5712–5719.
- [7] A. Pierson, W. Schwarting, S. Karaman, and D. Rus, “Learning Risk Level Set Parameters from Data Sets for Safer Driving,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, Paris, France: IEEE, Jun. 2019, pp. 273–280.
- [8] H. U. Ahmed, Y. Huang, and P. Lu, “A Review of Car-Following Models and Modeling Tools for Human and Autonomous-Ready Driving Behaviors in Micro-Simulation,” en, *Smart Cities*, vol. 4, no. 1, pp. 314–335, Mar. 2021.
- [9] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, “Planning for Autonomous Cars that Leverage Effects on Human Actions,” in *Robotics: Science and Systems XII*, Robotics: Science and Systems Foundation, 2016.
- [10] D. Sadigh, S. S. Sastry, and S. A. Seshia, “Verifying Robustness of Human-Aware Autonomous Cars,” en, *IFAC-PapersOnLine*, vol. 51, no. 34, pp. 131–138, 2019.
- [11] Z. Cao, E. Biyik, W. Z. Wang, et al., *Reinforcement Learning based Control of Imitative Policies for Near-Accident Driving*, arXiv:2007.00178 [cs, eess, stat], Jun. 2020.
- [12] A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer, “Imitating driver behavior with generative adversarial networks,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, USA: IEEE, Jun. 2017, pp. 204–211.
- [13] F. Codevilla, E. Santana, A. M. López, and A. Gaidon, *Exploring the Limitations of Behavior Cloning for Autonomous Driving*, arXiv:1904.08980 [cs], Apr. 2019.
- [14] O. Siebinga, A. Zgonnikov, and D. Abbink, “A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles,” en, *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 4, pp. 1–21, Dec. 2022.
- [15] Y. Lyu, W. Luo, and J. M. Dolan, “Adaptive safe merging control for heterogeneous autonomous vehicles using parametric control barrier functions,” in *2022 IEEE Intelligent Vehicles Symposium (IV)*, IEEE, 2022, pp. 542–547.
- [16] G. Li, S. E. Li, B. Cheng, and P. Green, “Estimation of driving style in naturalistic highway traffic using maneuver transition probabilities,” en, *Transportation Research Part C: Emerging Technologies*, vol. 74, pp. 113–125, Jan. 2017.
- [17] J. Wang, Y. Zheng, X. Li, C. Yu, K. Kodaka, and K. Li, “Driving risk assessment using near-crash database through data mining of tree-based model,” en, *Accident Analysis & Prevention*, vol. 84, pp. 54–64, Nov. 2015.
- [18] T. A. Dingus, F. Guo, S. Lee, et al., “Driver crash risk factors and prevalence evaluation using naturalistic driving data,” en, *Proceedings of the National Academy of Sciences*, vol. 113, no. 10, pp. 2636–2641, Mar. 2016.
- [19] Y. Lyu, W. Luo, and J. M. Dolan, “Risk-aware safe control for decentralized multi-agent systems via dynamic responsibility allocation,” in *2023 IEEE International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2023.
- [20] Y. Lyu, J. M. Dolan, and W. Luo, “Decentralized safe navigation for multi-agent systems via risk-aware weighted buffered voronoi cells,” in *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 2023, pp. 1476–1484.
- [21] Y. Lyu, J. M. Dolan, and W. Luo, “Cbf-inspired weighted buffered voronoi cells for distributed multi-agent collision avoidance,” in *2023 American Control Conference (ACC)*, IEEE, 2023, pp. 4513–4518.
- [22] M. Montanino and V. Punzo, “Making NGSIM Data Usable for Studies on Traffic Flow Theory: Multistep Method for Vehicle Trajectory Reconstruction,” en, *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2390, no. 1, pp. 99–111, Jan. 2013.
- [23] J. Morton, T. A. Wheeler, and M. J. Kochenderfer, “Analysis of Recurrent Neural Networks for Probabilistic Modeling of Driver Behavior,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 5, pp. 1289–1298, May 2017.
- [24] C. Thiemann, M. Treiber, and A. Kesting, “Estimating Acceleration and Lane-Changing Dynamics from Next Generation Simulation Trajectory Data,” en, *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2088, no. 1, pp. 90–101, Jan. 2008.
- [25] N. Deo and M. M. Trivedi, *Convolutional Social Pooling for Vehicle Trajectory Prediction*, arXiv:1805.06771 [cs], May 2018.
- [26] Y. Zhang, S. Zhang, and R. Luo, “Lane Change Intent Prediction Based on Multi-Channel CNN Considering Vehicle Time-Series Trajectory,” in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, Macau, China: IEEE, Oct. 2022, pp. 287–292.
- [27] S. Su, K. Muelling, J. M. Dolan, P. Palanisamy, and P. Mudalige, “Learning vehicle cooperative lane-changing behavior from observed trajectories in the ngsim dataset,” in *Proceedings of IEEE Intelligent Vehicles Symposium (IV '18)*, Jun. 2018, pp. 1412–1417.
- [28] Y. Lyu, W. Luo, and J. M. Dolan, “Probabilistic safety-assured adaptive merging control for autonomous vehicles,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021, pp. 10 764–10 770.
- [29] Y. Lyu, W. Luo, and J. M. Dolan, “Responsibility-associated multi-agent collision avoidance with social preferences,” in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, IEEE, 2022, pp. 3645–3651.
- [30] S. Wang, Y. Lyu, and J. M. Dolan, *Active Probing and Influencing Human Behaviors Via Autonomous Agents*, arXiv:2304.11817 [cs], Apr. 2023.