
Injecting Knowledge in Data-driven Vehicle Trajectory Predictors

Mohammadhossein Bahari

Ismail Nejjar

Alexandre Alahi

Ecole polytechnique fédérale de lausanne (EPFL)
mohammadhossein.bahari@epfl.ch

Abstract

Vehicle trajectory prediction tasks have been commonly tackled from two distinct perspectives: either with knowledge-driven methods or more recently with data-driven ones. On the one hand, we can explicitly implement domain-knowledge or physical priors such as anticipating that vehicles will follow the middle of the roads. While this perspective leads to feasible outputs, it has limited performance due to the difficulty to hand-craft complex interactions in urban environments. On the other hand, recent works use data-driven approaches which can learn complex interactions from the data leading to superior performance. However, generalization, *i.e.*, having accurate predictions on unseen data, is an issue leading to unrealistic outputs. In this paper, we propose to learn a "Realistic Residual Block" (RRB), which effectively connects these two perspectives. Our RRB takes any off-the-shelf knowledge-driven model and finds the required residuals to add to the knowledge-aware trajectory. Our proposed method outputs realistic predictions by confining the residual range and taking into account its uncertainty. We also constrain our output with Model Predictive Control (MPC) to satisfy kinematic constraints. Using a publicly available dataset, we show that our method outperforms previous works in terms of accuracy and generalization to new scenes. We will release our code and data split here: <https://github.com/vita-epfl/RRB>.

1 Introduction

While driving, humans have this powerful capability to anticipate other drivers' decisions. Similarly, an autonomous vehicle should have the same prediction capability to safely navigate alongside human drivers. Some researchers addressed the vehicle trajectory prediction task, also known as microscopic traffic modeling, by building hand-crafted functions based on the available domain-knowledge to model average driving behaviors [28, 11, 52, 10, 9, 56]. These methods are interpretable and usually lead to a set of feasible predictions. However, they have limited performance since they not only miss non-average behaviors, but also are not able to model complex interactions. Conversely, recent works solely rely on experience, *i.e.*, learning from data, mostly using neural networks [59, 55, 1, 2, 32, 50]. Using large amount of data helps these methods to achieve accurate predictions without explicitly modeling the domain-knowledge. However, their predictions are not essentially realistic (on-road) and in some cases, even counter-intuitive. Moreover, they are prone to overfitting on the training data or poor performance on out-of-distribution data. Combining the domain-knowledge and the data will benefit from the strengths of both approaches and avoid their shortcomings.

We argue that driving is a skill learned from domain-knowledge and experience. The former is typically driven by physical constraints such as respecting road constraints or avoiding collisions. The later is driven by social conventions *e.g.*, the way drivers interact with each other or the safe way to enter a roundabout. Consequently, learning to predict vehicle trajectories can be re-framed as

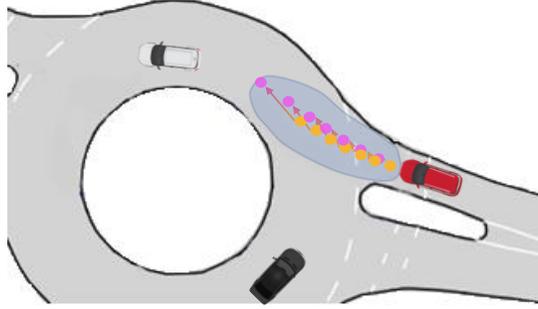


Figure 1: Illustration of our Realistic Residual Block (RRB) model. The knowledge-based model generates a scene-compliant trajectory (yellow dots) which could not effectively account for other agents, hence, is too conservative. Our data-driven RRB improves the prediction (purple dots) by adding confined residuals conditioned on other agents in the scene (illustrated by the arrows). The blue region shows the physically-constrained output space for our RRB predictions.

learning to combine both knowledge-driven and data-driven methods. Human’s ability to employ both knowledge and data is not specific to driving. Researchers believe that humans learn rich representations (we refer to as knowledge) as well as patterns from observed examples in everyday life [34]. This allows them to learn from fewer examples and to generalize to new conditions [33].

A popular way of adding domain-knowledge to the neural network is by adding constraints to the problem and optimizing the network under these constraints. Authors in [42] proposed a road-loss which can be interpreted as an approximation of the scene constraint in order to avoid off-road predictions. However, directly optimizing the model under constraints makes the optimization difficult and leads to sub-optimal results [17, 44].

There exist two main challenges towards creating a knowledge-driven and data-driven model. First, the combined model should be differentiable so that the data-driven part can be trained. This means that the integration needs careful design as the knowledge-driven part of the model usually is not differentiable. The second challenge is to preserve the benefits of both worlds after merging them *i.e.*, the final output should be realistic despite the fact that neural network might generate unrealistic outputs.

In this work, we address the aforementioned challenges. In the proposed approach shown in figure 1, the knowledge-driven (KD) trajectory is achieved by any knowledge-driven model. The KD output is then taken as input by our Realistic Residual Block (RRB) and the residuals required to be added to KD trajectory are found. In other words, the KD trajectory estimates coarse-grained behaviors based on the common driving performances while the residuals address fine-grained behaviors coming from non-modeled social interactions in KD trajectory as well as the long-tail of performances. This structure allows imposing knowledge by any function while allowing the residual block to be trainable. Moreover, in order not to diminish the feasibility of KD prediction, we can physically constrain the output of our RRB (*i.e.*, the allowed feasible offset) and leverage its uncertainty in the combination. Hence, our RRB uses a physically-constrained Inverse-variance weighting approach to add feasible and confident residuals to the KD trajectory. We finally constrain the final output with Model Predictive Controller (MPC) to ensure kinematic-feasibility of predictions.

The contributions of this work can be summarized as: (1) proposing a Realistic Residual Block (RRB) as an effective way of leveraging both knowledge and data in vehicle trajectory prediction. Our RRB complements the knowledge-driven output with realistic and confident outputs. (2) Using Model Predictive Controller (MPC) to bring kinematic constraints to the vehicle trajectory prediction task with latent control variables. (3) Demonstrate the generalization of our approach to new scenes and the multimodal case.

2 Related Work

Vehicle trajectory prediction: Pioneering works addressed vehicle trajectory prediction problem by means of knowledge-driven methods. Researchers in [11] used Kalman filter [24] to predict vehicle

future trajectory. In order to incorporate scene information, [60] proposed associating vehicle’s positions with the lanes of the road. Vehicle-vehicle interaction is addressed in [52] to predict the longitudinal motion of a target vehicle. In [56], an asymmetric optimal velocity model is presented to capture the asymmetry between acceleration and deceleration and [28] studied lane change decision empirically. On the other hand, many researchers tackle the problem by leveraging data-driven models. A deep belief network is proposed in [55] to model lane-changing behavior. Authors in [59, 14] model lane-changing and car following behaviors simultaneously by employing long short-term memory (LSTM) neural networks and convolutional social pooling respectively. Authors in [47] used an attention module to incorporate scene features into an LSTM model. An inverse optimal control (IOC) ranking module is used in [36] to determine the most likely hypotheses incorporating scene context and interactions. Moreover, some researchers addressed the multimodal nature of human behavior prediction by using Winner-Takes-All (WTA) loss method [13, 41]. While all mentioned works rely on knowledge or data exclusively, our solution benefits from both approaches.

Injecting domain-knowledge to neural networks: Knowledge has been incorporated into the non-data-driven models to preserve realism and improve generalization. In [3, 51], space-time constraints were employed to confine the outputs to the feasible range. Researchers have proposed different means for incorporating domain-knowledge into neural networks [54, 6, 45]. It can be injected into neural networks by designing specialized network structures [16, 18]. As an example, Convolutional neural network (CNN) was created by changing multilayer perceptron (MLP) networks, considering image specifications [35]. Another approach for fusing neural networks with domain knowledge is to modify the learning algorithm. This can be done by exposing knowledge-driven negative samples to the model [40] or adding constraints to the outputs of the network [45]. Other approaches are adding knowledge to the training data [26] and post-processing [15]. From another perspective, some researchers used neural networks to address the imperfection of knowledge-driven models. They used knowledge-driven models to find biases of the data and the neural network compensates the error by outputting required residuals. [57, 49] predict residuals on top of a physics-based robotic controller and [25] finds residual minimiser of numerically discretized differential equations.

A key limitation in residual modeling is by adding neural network-based residuals, the physics-based constraints, which are required for a realistic prediction, are hard to preserve [54]. In this paper, we mitigate this limitation by confining the residual values. Moreover, we replaced naive addition of residuals by variance-based integration methods.

In the context of vehicle trajectory prediction, previous works attempt to add scene knowledge by adding penalty terms to the loss function. Researchers in [5] defined an on-road loss to keep the predictions inside the road. The proposed loss requires the output to be an occupancy heatmap, which is computationally expensive to achieve and also not compatible with most of the trajectory prediction works [42]. Authors in [42] proposed an off-road loss defined as the euclidean distance between each predicted waypoint and the nearest drivable point to penalize off-road predictions. Their experiments show that although off-road predictions are reduced, the performance deteriorates in terms of ℓ_2 loss. This is due to the fact that direct optimization of the model with constraints which are non-linear with respect to model parameters is difficult and leads to sub-optimal solutions [17, 44]. In this paper, we propose a new method for adding scene knowledge to the model.

Kinematically-feasible predictions: Kinematic constraints are physical rules that need to be satisfied for a realistic vehicle prediction. Authors in [12] showed that the predictions of the neural network model were not essentially kinematically-possible. Researchers in [12, 38] solved the problem by adding the bicycle model [30] as a layer to the model. The network estimates control commands instead of coordinates and the kinematic layer converts them to a feasible trajectory. Despite the effectiveness of this approach, in many cases, it cannot be employed as most of the off-the-shelf models predict coordinates rather than control commands. Moreover, having control commands as the outputs prevents any further knowledge injection to the model’s output, since any changes to the predicted trajectory can lead to a non-kinematically-feasible trajectory. In contrast, we propose using a Model Predictive Controller (MPC) to satisfy kinematic constraints. This allows the outputs to be in the coordinate format while the control commands are latent variables and mitigates the mentioned problems. Model Predictive Controller is commonly used for planning [60, 7, 27, 23, 4]. In this work, we show how it can also be used for the vehicle trajectory prediction task.

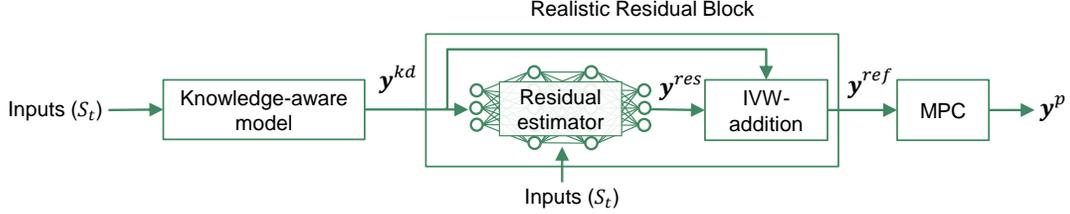


Figure 2: Our proposed Realistic Residual Block (RRB) takes as input (i) the state of the scene, S_t , and (ii) the knowledge-driven prediction, y^{kd} . The Residual estimator block builds physically-constrained residuals y^{res} . Then, the IVW-addition block merges y^{kd} and y^{res} by Inverse-Variance Weighted sum and forms y^{ref} according to equation 3. Finally, Model Predictive Controller (MPC) satisfies the kinematic constraints.

3 Proposed Method

Humans have a clear understanding of the domain-knowledge while driving *e.g.*, where the drivable and non-drivable areas are. Moreover, they learn specific aspects of driving by experience *e.g.*, interacting with other agents. However, benefiting from the domain-knowledge and learning from experience simultaneously is challenging. We address the problem by proposing a Realistic Residual Block (RRB). RRB finds data-driven residuals conditioned on the knowledge-aware prediction. The output is a physically-constrained Inverse-Variance Weighted (IVW) sum of the knowledge-driven (KD) trajectory with the residuals. Finally, Model Predictive Control (MPC) is incorporated to satisfy kinematic constraints. Figure 2 shows a high level picture of our proposed model. We will explain each part of the model in the following subsections.

3.1 Problem Formulation

The goal of our model is to predict future positions of a vehicle given its history and surroundings. Therefore, the state of the scene input to the model at time t , S_t , consists of an image of the scene and histories of ego-vehicle and other vehicles in the scene. By ego-vehicle, we refer to the vehicle whose future is predicted. The observation vectors for ego-vehicle and other vehicles' history comprises of T_o observed 2D coordinates of them in the world space. Without loss of generality, each time, one of the agents in the scene is taken as the ego-vehicle. At time t , the model predicts y_t^{ref} for the next T_p prediction frames as a Gaussian distribution with mean μ and variance σ as:

$$p(\mathbf{y}_t^{ref} | S_t, \boldsymbol{\mu}_t, \boldsymbol{\sigma}_t) = \prod_{j=1}^{T_p} \mathcal{N}(y_{t+j}^{ref} | S_t, \mu_{t+j}, \sigma_{t+j}). \quad (1)$$

Note that we indicate the sequence of values in a bold text. We omit the index t in the rest of the paper for simplicity.

3.2 Knowledge-aware Prediction

Our proposed method is flexible in the choice of the knowledge-aware model. Hence, any off-the-shelf knowledge-driven prediction model can be employed to create the KD trajectory. In section 4.5, we show the robustness of our RRB to different knowledge-driven models. In this work, we make the KD trajectory by utilizing the scene knowledge, the most influential domain-knowledge in vehicle prediction task. Inspired by [60], we use the lanes of the road to form the scene-compliant trajectory y^{kd} . To have a probabilistic framework, we take KD trajectory as the mean of a Gaussian distribution with fixed variances which are approximated by the statistics of the training data.

3.3 Realistic Residual Block

While KD trajectory reflects average behaviors with respect to the scene, our proposed data-driven Realistic Residual Block learns the missing complex interactions. The architecture of RRB is depicted in Figure 3. The learned residuals are first confined and then merged with the KD prediction by the

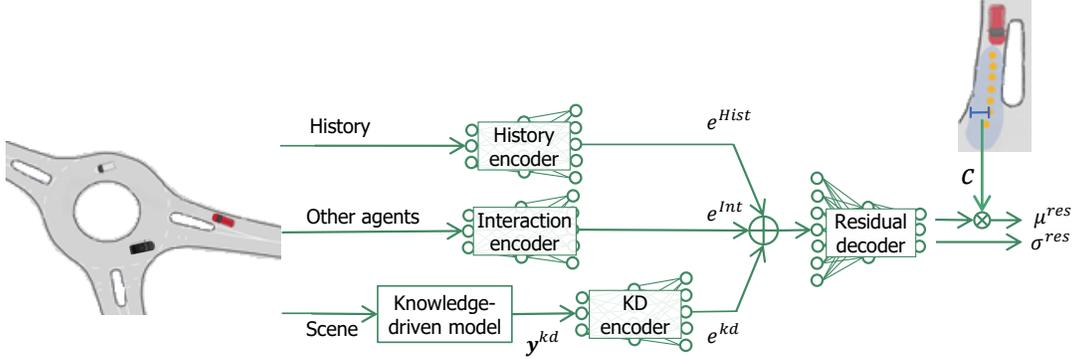


Figure 3: The 'Residual estimator' block inside our RRB. Three encoders are employed to embed the history, the interactions and the KD trajectory. The residual decoder integrates the features and estimates the residual distribution. The mean of the distribution is confined to C , a real-word extracted parameter.

IVW-addition block. This will lead to feasible and confident predictions. Both Residual estimator and integration mechanism are explained in the following subsections.

3.3.1 Residual Estimator

Our Residual estimator's structure is shown in Figure 3. The inputs are the KD trajectory and S_t consisting of the image of the scene and history of agents in coordinate format. The history of ego-vehicle is processed by the history encoder to find the driver's intention leading to the feature e^{Hist} . We model the interaction among agents in a simple yet effective approach. Inspired by how vehicles interact in the real-world, our preprocessing eliminates non-interacting agents in the scene. In contrast with most previous works [1, 50] which consider surrounding agents, it excludes all the agents behind the vehicle as usually, vehicles interact with the agents in front. Moreover, as only the closest vehicles can impact the driver's behavior, it only preserves a set of closest vehicles and passes them to the interaction encoder model to get the interaction features e^{Int} . Our experiments show the effectiveness of the model in learning interactions between agents. The KD trajectory y^{kd} , is encoded into the vector e^{kd} by the KD encoder network. All the encoders are feed-forward networks with ReLU non-linearities. Finally, the residual decoder estimates the residuals given the concatenated features as:

$$\begin{aligned} (\hat{\mu}^{res}, \sigma^{res}) &= MLP([e^{Hist}, e^{Int}, e^{kd}]; W^{res}), \\ \mu^{res} &= C \hat{\mu}^{res}, \end{aligned} \quad (2)$$

where $(\mu^{res}, \sigma^{res})$ are parameters of the Gaussian distribution y^{res} and MLP is a multilayer perceptron network parameterized by W^{res} . Note that we bound the range of $\hat{\mu}^{res}$ to $(-1, 1)$ by using tanh activation function in the last layer. μ^{res} is the scaled version of $\hat{\mu}^{res}$ by C to adjust the max feasible deviation from the center of the road. C is a real-word extracted parameter which can be adapted to each scene. In our experiments, we chose C equal to the half of the minimum width of the road in each scene.

3.3.2 IVW-addition

The merged trajectory y^{ref} can be achieved by simply adding residual Gaussians to the KD Gaussian prediction, referred to "A-RRB" baseline in Section 4.3. However, adding uncertain residuals to the KD trajectory can impair the KD predictions. To mitigate that, we utilize Inverse-Variance Weighting (IVW) [8] to scale residuals according to their uncertainties which results in the most certain output (in terms of output variance) among all weighted averages [21]. Hence, our goal is to find the weights w, \tilde{w} such that each point on the merged trajectory shown in Equation 3 has the minimum variance.

$$y^{ref} = w y^{kd} + \tilde{w} (y^{kd} + y^{res}). \quad (3)$$

Let's represent $y^{kd} + y^{res}$ by y^{ad} for brevity. The points on each trajectory are assumed temporally-independent. Weight matrices have the form of $w = \text{diag}(w_1, w_2)$, $\tilde{w} = \text{diag}(\tilde{w}_1, \tilde{w}_2)$ where w_i, \tilde{w}_i

are scalar values. Note that \mathbf{y}^{kd} is a prior for \mathbf{y}^{res} and hence they are not independent. The problem formulation is as follows:

$$\begin{aligned} & \underset{w, \tilde{w}}{\operatorname{argmin}} (\sigma_{11}^{ref}, \sigma_{22}^{ref}), \\ & \text{subject to: } \sigma_{ref} = w\sigma^{kd}w^T + \tilde{w}\sigma^{ad}\tilde{w}^T + 2w\sigma^{kd,ad}\tilde{w}^T, \\ & \quad w + \tilde{w} = I_{2 \times 2}, \end{aligned} \quad (4)$$

where σ^{kd}, σ^{ad} are the covariance matrices and $\sigma^{kd,ad}$ is the cross-covariance matrix. Solving the constraint problem using Lagrangian multipliers leads to the following solution:

$$w = \begin{bmatrix} \frac{\sigma_{11}^{ad} - \sigma_{11}^{kd,ad}}{\sigma_{11}^{ad} + \sigma_{11}^{kd} - 2\sigma_{11}^{kd,ad}} \\ \frac{\sigma_{22}^{ad} - \sigma_{22}^{kd,ad}}{\sigma_{22}^{ad} + \sigma_{22}^{kd} - 2\sigma_{22}^{kd,ad}} \end{bmatrix}, \tilde{w} = \begin{bmatrix} \frac{\sigma_{11}^{kd} - \sigma_{11}^{kd,ad}}{\sigma_{11}^{ad} + \sigma_{11}^{kd} - 2\sigma_{11}^{kd,ad}} \\ \frac{\sigma_{22}^{kd} - \sigma_{22}^{kd,ad}}{\sigma_{22}^{ad} + \sigma_{22}^{kd} - 2\sigma_{22}^{kd,ad}} \end{bmatrix}. \quad (5)$$

3.4 Multimodal Prediction and Loss Function

Our approach can easily be extended to a multimodal predictor using Winner-Takes-All (WTA) loss (also known as oracle loss) approach [46]. To do so, the knowledge-aware model should generate multiple plausible future trajectories. RRB takes the KD predictions and finds the associated required residuals. Finally, the loss will be calculated, considering the closest mode to the ground truth. Minimizing the log-likelihood will lead to the following loss function:

$$l(\theta) = - \sum_{n=1}^N \sum_{m=1}^M \mathbb{1}(m = m^*) [\log p(\mathbf{x}_{n,m}^p | S_n, \mu_{n,m}^{ref}(\theta), \sigma_{n,m}^{ref}(\theta))], \quad (6)$$

where N and M are number of samples and modes respectively, $\mathbb{1}$ is the indicator function, and m^* is the closest output to the ground truth in terms of ℓ_2 distance. Note that for the single modal case, the equation holds with $M = 1$.

3.5 Model Predictive Control

To add kinematic feasibility, in contrast with the previous works [12, 38] which utilize a kinematic layer after computing control commands, we employ Model Predictive Control (MPC). This gives the model the flexibility to estimate the positions instead of finding the control commands which is beneficial in adding the domain-knowledge. MPC minimizes its cost function subject to a set of constraints. The state parameters for agent i in time t is $s_t = [x_t, y_t, \phi_t, v_t]$ which consists of coordinates, orientation and speed. We denote control parameters acceleration and the steering angle by $u_t = [a_t, \gamma_t]$. The dynamics of the system is formulated using bicycle model F_{bic} [30] which is shown to be sufficient for normal manoeuvres [31]. Then, the MPC solves the following optimization problem:

$$\begin{aligned} \mathbf{y}_t^p = & \underset{s_{t:t+T_p}, u_{t:t+T_p}}{\operatorname{argmin}} \sum_{j=1}^{T_p} \|s_{t+j}[0:1] - \mathbf{y}_{t+j}^{ref}\|_2^2 + \lambda \|u_{t+j} - u_{t+j-1}\|_2^2 \\ & \text{subject to: } s_{t+1} = F_{bic}(s_t, u_t), \quad s_0 = s_{init}, \quad u_{min} < u_t < u_{max}, \end{aligned} \quad (7)$$

where \mathbf{y}^{ref} is the reference trajectory, λ is a hyper-parameter and u_{min}, u_{max} are minimum and maximum feasible control values respectively.

4 Experiments

We evaluate our RRB against other baselines to test the following hypotheses: (1) RRB brings the advantages of both knowledge-driven and data-driven models thus, outperforms both types of models, (2) RRB structure merges knowledge-driven and data-driven models better than other fusion techniques, (3) adding MPC makes outputs more realistic by satisfying kinematic constraints, (4) our trained RRB can improve performance of different knowledge-driven models without the need for fine-tuning, (5) IVW-addition and residual confinement are essential parts for more realistic outputs.

4.1 Dataset

We evaluate the proposed method on Interaction dataset [58]. It is a large-scale real-world dataset which consists of top-down scenes from intersections, highways, and roundabouts. The data is collected from three different continents (North America, Asia and Europe). It includes locations of dynamic agents such as vehicles and pedestrians for each frame of the 10 Hz downsampled video. The dataset also provides the static context of the scene. Interaction dataset is challenging as it includes interactions between vehicles, different environments, and potentially multiple plausible predictions. We used the same settings as [58]. The observation and prediction lengths indicate the number of frames used to represent the past states and to be predicted, respectively. We have set observation lengths equal to 5 frames (2.5 s) and prediction length as 10 frames (5 s). The dataset includes scenes with different numbers of samples. Hence, in order not to be overwhelmed by scenes with large number of samples, we report the average of performance of a model on the three categories regardless of the number of samples in each category. Moreover, to study the model’s generalization power, we consider two scenarios: the first scenario is scene-overfitting, in which all scenes exist in the training set, but 20% of data is kept for the test set. In this scenario the model should be able to overfit on the scene and generalize on the interactions. We argue that to assess model generalization in terms of scene perception, another scenario named scene-generalization is required. In this scenario, the three following scenes, 'DR_USA_Intersection_MA', 'DR_USA_Roundabout_SR', 'DR_CHN_Merging_ZS' are kept for the test set and the rest scenes are used in training. This challenging scenario reveals the performance of the models in a new environment.

4.2 Implementation Details

The models are trained for 50 epochs with batch size of 32. We employed Adam optimizer [29] with the initial learning rate of 0.001, which is decreased by half every 10 epochs. The model is implemented using PyTorch [43]. Network’s building blocks are MLP networks with the following hidden layers: History and interaction encoders have (32,32,64), KD encoder has (32,64) and the decoder has (256,128,128,64) hidden layers.

4.3 Metrics and Baselines

The following metrics were used for the evaluation:

1. Average/final displacement error (**ADE/FDE**). Average displacement error (ADE) and Final displacement error (FDE) are adopted as two common evaluation metrics. In the multimodal case, similar to the previous works [20, 37], the closest mode to the ground truth is chosen.
2. Road violation (**RV**). Inspired by [42], we define this metric as the percentage of average number of points predicted in the off-road area. For the multimodal case, we average RVs for the modes weighted by their probability. RV measures the feasibility of predictions concerning the scene.
3. Cross track (**CT**). Cross track is the distance between the actual destination and the final point on the retimed predicted sequence by the ground truth speed profile [19]. In other words, cross track is the distance between the actual destination and the final point on the retimed predicted sequence. Cross track metric is able to express the spatial effectiveness of the model, excluding temporal aspects.

To demonstrate the effectiveness of RRB, we compare the results with the following baselines:

Naive baselines: We used Kalman filter as a linear prediction (**Lin**). We also report Constant velocity (**CV**) [48] as another naive baseline.

Knowledge-driven baselines: In order to form the KD trajectory explained in section 3.2, we utilize different strategies to predict vehicle velocity. We name Constant velocity strategy [48] as **KD1**, and Leader follower [52] which tackles interaction between agents as **KD2**.

Data-driven baselines: For the data-driven models that address agent-agent interactions, we report Social LSTM (**S-LSTM**) [1], Social Attention (**S-ATT**) [53] (numbers reported from [39]), and Social GAN (**S-GAN_M**) [20]. We report Social WaGDAT (**S-WaGDAT_M**) [39] (numbers reported from their paper) as a state-of-the-art work that models both scene and interactions. We also report the performance of our encoder-decoder neural network (**EDN**) similar to Figure 3 trained with equation

Table 1: Quantitative results of baselines in scene-generalization scenario. ADE/FDE and CT are reported in meters. The lower the better for all metrics.

Models		Scene-generalization		
		ADE/FDE	RV	CT
Naive	Lin	4.13 / 8.77	34	3.6
	CV [48]	3.12 / 7.34	24	3.19
Knowledge-driven	KD1 [48, 60]	2.92 / 6.62	0	1.87
	KD2 [52, 60]	2.85 / 6.55	0	1.74
Data-driven	S-LSTM [1]	2.85 / 7.17	72	4.26
	S-GAN_M [20]	2.34 / 5.82	52	5.71
	EDN	2.78 / 6.7	7	3.4
Mixed approaches	VI1	2.52 / 6.27	3	2.9
	VI2	2.62 / 6.38	2	2.32
	RL [42]	2.56 / 6.40	9	3.15
	RRB	2.44 / 6.04	0	1.98
	RRB_M	2.15 / 5.08	0	1.81
	RRB_M+MPC	2.13 / 5.02	0	1.81

Table 2: Quantitative results of baselines in scene-overfitting scenario. ADE/FDE are in meters.

Models	ADE/FDE
CV [48]	2.80 / 6.59
KD1 [48, 60]	2.59 / 6.00
KD2 [52, 60]	2.53 / 5.83
S-LSTM [1]	2.33 / 4.52
EDN	1.81 / 4.05
S-ATT [53]	2.29 / 4.25
S-GAN_M [20]	2.12 / 4.20
S-WaGDAT_M [39]	1.62 / 3.35
RRB_M	1.49 / 3.68

6 as the loss function. We replaced the knowledge-driven model in Figure 3 with a convolutional neural network to represent a fully data-driven model. Note that the multimodal baselines’ names end with ‘_M’.

Mixed baselines: We implemented the road loss (**RL**) [42] to impose the scene constraint to the model. It adds the road loss to the NN model to help it learn the drivable and non-drivable regions. We also consider two alternatives to the RRB model. The first one is to use a Variance-based Integration (**VI1**) between the KD1 and data-driven prediction EDN by employing IVW method [8]. Note that this is equivalent to using Bayes rule while KD output is taken as the prior and EDN as the likelihood. The second approach is to replace the KD variance in VI1 by a fixed hyperparameter (**VI2**).

Our proposed solution is RRB model (**RRB**) explained in section 3.3 which uses KD1 method. We express the effectiveness of our method on other KD predictions in the ablation study. We also demonstrate the performance of our model in the multimodal case (**RRB_M**) with 2 modes and when constrained by the MPC (**RRB_M+MPC**). Moreover, in the ablation study, we report the performance of RRB while IVW-addition is replaced by simple addition (**A-RRB**) and also a non-confined RRB model (**NC-RRB**).

4.4 Results

Table 1 provides the results in scene-generalization scenario. Knowledge-driven models outperform other models in terms of RV and CT metrics since they can perceive scene well and also generalize to new environment. The data-driven approaches have better ADE/FDE as they can learn interactions from data. Although EDN performs well in scene-overfitting scenario (shown in Table 2), it fails

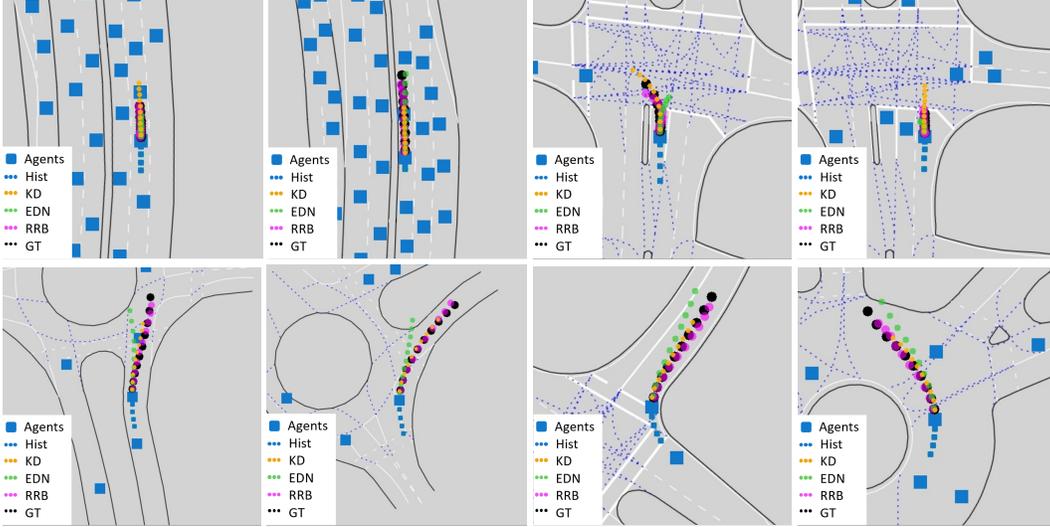


Figure 4: Qualitative results of different baselines. Given the scene and history (Hist), the models predict the future positions. The ground truth (GT) is shown in black. The encoder-decoder neural network (EDN) captures the interactions while is prone to create unrealistic outputs shown in first row. Knowledge-driven (KD) model has realistic predictions as it uses the scene knowledge but cannot reason about the interactions among vehicles as shown in the second row. Our proposed RRB predicts realistic and interaction-aware outputs.

in generalizing to the new environment leading to high RV and CT values. The large gap between the performance of data-driven methods in scene-overfitting and the scene-generalization scenarios reveals the limited generalization power of solely data-driven approaches, which can be mitigated by leveraging domain-knowledge. The proposed RRB model outperforms knowledge-driven models as it learns vehicle-vehicle interactions from data and goes beyond average behavior. Compared to the data-driven models, RRB can better generalize to new environments due to utilization of the scene knowledge. Compared with the mixed approaches, RRB outperforms both VI1 and VI2 because as opposed to them that merge two independent models, RRB generates residuals conditioned on KD output. Our method has zero RV because of the realistic residual scale which prevents off-road predictions. Note that RL could not improve EDM performance as it is very sensitive to the choice of hyperparameters and has many local minima.

The possibility of having multiple outputs is investigated by increasing the number of modes of RRB model to 2. The results approve that RRB_M is able to successfully capture multiple modes while it provides the required residuals for each mode individually leading to a multimodal knowledge-aware model. MPC is employed in RRB_M+MPC in order to ensure the kinematic-feasibility of the outputs. The results show that kinematic constraints are satisfied without loss of performance. Hence, the model better approximates human behavior.

The results of previous works in scene-overfitting scenario are reported in Table 2. Our RRB_M model outperforms previous works and especially the recent S-WaGDAT_M [39] model which expresses the ability of the model to learn from data. Note that RRB_M could successfully improve performance of KD1 by adding the residuals.

We visualize the outputs of different models in Figure 4. The first row images visualize the cases where KD prediction is not accurate because of the missing interactions. RRB complements KD trajectory by accounting for interactions in the residuals. The second row images emphasize on realism of the predictions. Although EDN can reason about interactions, it can have unrealistic predictions with respect to the scene especially in a new environment. However, the realistic elements of RRB preserve the feasibility of the final output as shown in all images.

Table 3: Results of the two ablation studies on robustness and realistic elements of the proposed RRB

(a) Assessing the robustness of RRB with respect to different knowledge-driven models. RRB is not trained in this experiment but we used the fixed model in Table 1.

Models	Scene-generalization		
	ADE/FDE	RV	CT
Lin	4.13 / 8.77	34	3.6
Lin+RRB	3.11 / 7.50	2	3.59
KD2	2.85 / 6.55	0	1.74
KD2+RRB	2.49 / 6.14	0	2.04

(b) Ablation study on the two realistic elements of the proposed RRB, residual confinement and IVW-addition block.

Models	Scene-generalization		
	ADE/FDE	RV	CT
NC-RRB	2.28 / 5.59	10	1.94
A-RRB	2.47 / 6.08	3	2.07
RRB	2.44 / 6.04	0	1.98

4.5 Ablation Study

In this section, a set of ablation studies are performed to shed light on the effectiveness of different parts of the model. First, we want to study the robustness of RRB concerning the different knowledge-driven models. We report two knowledge-driven methods in Table 1 and report the RRB performance while KD1 was employed. Without retraining the model, we replaced KD1 with Lin (Lin+RRB) and KD2 (KD2+RRB) models. The results are expressed in Table 3a, where the performance of the knowledge-driven model is reminded in the parenthesis. It shows that RRB could successfully improve the performance of the models in all metrics. This experiment shows the robustness and effectiveness of the proposed RRB for different KD trajectories.

The second ablation study aims at assessing realistic elements of RRB which are the realistic parameter c and IVW-addition. We trained the model without the limitation on the range of RRB as non-confined RRB (NC-RRB). In addition, IVW-addition can be replaced with simply adding the residuals to the KD prediction which is reported as A-RRB. This is equivalent to setting $w = \text{diag}(0, 0)$, $\tilde{w} = \text{diag}(1, 1)$ in Equation 3. The results are expressed in Table 3b. Non-confined model has more freedom in improving the KD trajectory hence achieves a better ADE/FDE. But it is prone to unrealistic predictions with large RV error. Also, neglecting uncertainties in A-RRB will increase RV and CT as the network utilized uncertain residuals. These experiments approve that the added elements play important roles in having realistic predictions.

5 Conclusions and Future Work

In this paper, we addressed the safety-critical task of vehicle trajectory prediction (also known as microscopic traffic modeling). We argue that neither solely knowledge-driven nor purely data-driven models can provide generalizable and accurate predictions. Given a set of experiments on a real-world dataset, our experiments demonstrate that knowledge-driven models have inferior performance but can generalize to new scenes. On the other hand, although data-driven models provide more accurate predictions, they tend to have unrealistic predictions in new scenes. Our proposed solution, named RRB, effectively merges knowledge-driven with data-driven models by finding residuals required to be added to the knowledge-driven prediction in order to model human stochastic behavior. It leverages physically-constrained Inverse-variance weighting approach to build realistic and confident residuals. To further make the predictions realistic, we employed a Model Predictive Control (MPC) to bring kinematic constraints to the final output. RRB outperforms all other counterparts in accuracy and generalizability. As future work, we will study the effectiveness of RRB when employing more complex knowledge-driven methods on a variety of agents, including pedestrians. Pedestrians are not constrained to the roads but respect specific social rules. We can take the existing knowledge-driven models such as ‘‘Social force’’ [22] as prior predictions and learn residuals required to better capture pedestrians dynamics. Moreover, our approach can help in complex multi-agent environments where there exist vehicles, pedestrians and cyclists each one with its own constraints and dynamics. In such scenarios, RRB can benefit from using strong knowledge-driven priors for each category and provide safe and accurate data-driven residuals. We hope that our work will pave the way to more methods combining the best of knowledge and data driven approaches.

6 Acknowledgements

This project was funded by Honda R&D Co., Ltd and the European union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement N 754354. Also, we would like to thank the reviewers for their helpful comments.

References

- [1] Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social lstm: Human trajectory prediction in crowded spaces. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
- [2] Alahi, A., Ramanathan, V., Goel, K., Robicquet, A., Sadeghian, A.A., Fei-Fei, L., Savarese, S.: Learning to predict human behavior in crowded scenes. In: Group and Crowd Behavior for Computer Vision, pp. 183–207. Elsevier (2017)
- [3] An, L., Tsou, M.H., Crook, S.E., Chun, Y., Spitzberg, B., Gawron, J.M., Gupta, D.K.: Space-time analysis: Concepts, quantitative methods, and future directions. *Annals of the Association of American Geographers* **105**(5), 891–914 (2015)
- [4] Bae, S., Saxena, D., Nakhaei, A., Choi, C., Fujimura, K., Moura, S.J.: Cooperation-aware lane change maneuver in dense traffic based on model predictive control with recurrent neural network. arXiv [abs/1909.05665](https://arxiv.org/abs/1909.05665) (2019)
- [5] Bansal, M., Krizhevsky, A., Ogale, A.S.: Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. In: Robotics: Science and Systems XV. vol. 15 (2019)
- [6] Borghesi, A., Baldo, F., Milano, M.: Improving deep learning models via constraint-based domain knowledge: a brief survey. arXiv preprint [arXiv:2005.10691](https://arxiv.org/abs/2005.10691) (2020)
- [7] Chen, Y., Hu, C., Wang, J.: Human-centered trajectory tracking control for autonomous vehicles with driver cut-in behavior prediction. *IEEE Transactions on Vehicular Technology* **68**(9), 8461–8471 (2019)
- [8] Cochran, W.G.: The combination of estimates from different experiments. *Biometrics* **10**(1), 101–129 (1954)
- [9] Coscia, P., Castaldo, F., Palmieri, F.A., Alahi, A., Savarese, S., Ballan, L.: Long-term path prediction in urban scenarios using circular distributions. *Image and Vision Computing* **69**, 81–91 (2018)
- [10] Coscia, P., Castaldo, F., Palmieri, F.A., Ballan, L., Alahi, A., Savarese, S.: Point-based path prediction from polar histograms. In: 2016 19th International Conference on Information Fusion (FUSION). pp. 1961–1967. IEEE (2016)
- [11] Cosgun, A., Ma, L., Chiu, J., Huang, J., Demir, M., Anon, A.M., Lian, T., Tafish, H., Al-Stouhi, S.: Towards full automated drive in urban environments: A demonstration in gomentum station, california. In: 2017 IEEE Intelligent Vehicles Symposium (IV). pp. 1811–1818. IEEE (2017)
- [12] Cui, H., Nguyen, T., Chou, F.C., Lin, T.H., Schneider, J., Bradley, D., Djuric, N.: Deep kinematic models for physically realistic prediction of vehicle trajectories. arXiv preprint [arXiv:1908.00219](https://arxiv.org/abs/1908.00219) (2019)
- [13] Cui, H., Radosavljevic, V., Chou, F.C., Lin, T.H., Nguyen, T., Huang, T.K., Schneider, J., Djuric, N.: Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 2090–2096. IEEE (2019)
- [14] Deo, N., Trivedi, M.M.: Convolutional social pooling for vehicle trajectory prediction. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) pp. 1549–15498 (2018)
- [15] Fang, Y., Kuan, K., Lin, J., Tan, C., Chandrasekhar, V.: Object detection meets knowledge graphs (2017)
- [16] França, M.V., Zaverucha, G., Garcez, A.S.d.: Fast relational learning using bottom clause propositionalization with artificial neural networks. *Machine learning* **94**(1), 81–104 (2014)
- [17] Ganchev, K., Gillenwater, J., Taskar, B., et al.: Posterior regularization for structured latent variable models. *Journal of Machine Learning Research* **11**(Jul), 2001–2049 (2010)
- [18] Garcez, A.d., Gori, M., Lamb, L.C., Serafini, L., Spranger, M., Tran, S.N.: Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. arXiv preprint [arXiv:1905.06088](https://arxiv.org/abs/1905.06088) (2019)
- [19] Gong, C., McNally, D.: A methodology for automated trajectory prediction analysis. In: AIAA Guidance, Navigation, and Control Conference and Exhibit. p. 4788

- [20] Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., Alahi, A.: Social gan: Socially acceptable trajectories with generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2255–2264 (2018)
- [21] Hartung, J., Knapp, G., Sinha, B.K.: Statistical meta-analysis with applications, vol. 738. John Wiley & Sons (2011)
- [22] Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. *Physical Review E* **51** (05 1998). <https://doi.org/10.1103/PhysRevE.51.4282>
- [23] Jeong, Y., Kim, S., Yi, K.: Surround vehicle motion prediction using lstm-rnn for motion planning of autonomous vehicles at multi-lane turn intersections. *IEEE Open Journal of Intelligent Transportation Systems* **1**, 2–14 (2020). <https://doi.org/10.1109/OJITS.2020.2965969>
- [24] Kalman, R.E.: A new approach to linear filtering and prediction problems (1960)
- [25] Kani, J.N., Elsheikh, A.H.: Dr-rnn: A deep residual recurrent neural network for model reduction. arXiv preprint arXiv:1709.00939 (2017)
- [26] Karpatne, A., Watkins, W., Read, J., Kumar, V.: Physics-guided neural networks (pgnn): An application in lake temperature modeling. arXiv preprint arXiv:1710.11431 (2017)
- [27] Kazemi, H., Mahjoub, H.N., Tahmasbi-Sarvestani, A., Fallah, Y.P.: A learning-based stochastic mpc design for cooperative adaptive cruise control to handle interfering vehicles. *IEEE Transactions on Intelligent Vehicles* **3**(3), 266–275 (2018)
- [28] Keyvan-Ekbatani, M., Knoop, V.L., Daamen, W.: Categorization of the lane change decision process on freeways. *Transportation research part C: emerging technologies* **69**, 515–526 (2016)
- [29] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv:1412.6980 (2014)
- [30] Kong, J., Pfeiffer, M., Schildbach, G., Borrelli, F.: Kinematic and dynamic vehicle models for autonomous driving control design. In: 2015 IEEE Intelligent Vehicles Symposium (IV). pp. 1094–1099. IEEE (2015)
- [31] Kong, J., Pfeiffer, M., Schildbach, G., Borrelli, F.: Kinematic and dynamic vehicle models for autonomous driving control design. In: 2015 IEEE Intelligent Vehicles Symposium (IV). pp. 1094–1099. IEEE (2015)
- [32] Kothari, P., Kreiss, S., Alahi, A.: Human trajectory forecasting in crowds: A deep learning perspective. arXiv preprint arXiv:2007.03639 (2020)
- [33] Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction. *Science* **350**(6266), 1332–1338 (2015)
- [34] Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S.J.: Building machines that learn and think like people. *Behavioral and brain sciences* **40** (2017)
- [35] LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
- [36] Lee, N., Choi, W., Vernaza, P., Choy, C., H. S. Torr, P., Chandraker, M.: Desire: Distant future prediction in dynamic scenes with interacting agents. pp. 2165–2174 (07 2017). <https://doi.org/10.1109/CVPR.2017.233>
- [37] Lee, N., Choi, W., Vernaza, P., Choy, C.B., Torr, P.H., Chandraker, M.: Desire: Distant future prediction in dynamic scenes with interacting agents. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 336–345 (2017)
- [38] Li, J., Ma, H., Tomizuka, M.: Conditional generative neural system for probabilistic trajectory prediction. 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) pp. 6150–6156 (2019)
- [39] Li, J., Ma, H., Zhang, Z., Tomizuka, M.: Social-wagdat: Interaction-aware trajectory prediction via wasserstein graph double-attention network. arXiv preprint arXiv:2002.06241 (2020)
- [40] Liu, Y., Yan, Q., Alahi, A.: Social nce: Contrastive learning of socially-aware motion representations. arXiv preprint arXiv:2012.11717 (2020)
- [41] Makansi, O., Ilg, E., Cicek, O., Brox, T.: Overcoming limitations of mixture density networks: A sampling and fitting framework for multimodal future prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7144–7153 (2019)

- [42] Niedoba, M., Cui, H., Luo, K., Hegde, D., Chou, F.C., Djuric, N.: Improving movement prediction of traffic actors using off-road loss and bias mitigation. *Machine Learning for Autonomous Driving Workshop at the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*
- [43] Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017 NiPS Talk)
- [44] Pathak, D., Krahenbuhl, P., Darrell, T.: Constrained convolutional neural networks for weakly supervised segmentation. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1796–1804 (2015)
- [45] von Rueden, L., Mayer, S., Beckh, K., Georgiev, B., Giesselbach, S., Heese, R., Kirsch, B., Pfrommer, J., Pick, A., Ramamurthy, R., et al.: Informed machine learning—a taxonomy and survey of integrating knowledge into learning systems. *arXiv preprint arXiv:1903.12394* (2020)
- [46] Rupprecht, C., Laina, I., DiPietro, R., Baust, M., Tombari, F., Navab, N., Hager, G.D.: Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3591–3600 (2017)
- [47] Sadeghian, A., Legros, F., Voisin, M., Vesel, R., Alahi, A., Savarese, S.: Car-net: Clairvoyant attentive recurrent network. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018*. pp. 162–180. Springer International Publishing, Cham (2018)
- [48] Schöller, C., Aravantinos, V., Lay, F., Knoll, A.: What the constant velocity model can teach us about pedestrian motion prediction. *IEEE Robotics and Automation Letters* **5**, 1696–1703 (2020)
- [49] Silver, T., Allen, K., Tenenbaum, J., Kaelbling, L.: Residual policy learning. *arXiv preprint arXiv:1812.06298* (2018)
- [50] Tang, C., Salakhutdinov, R.R.: Multiple futures prediction. In: *Advances in Neural Information Processing Systems*. pp. 15398–15408 (2019)
- [51] Tang, J., Song, Y., Miller, H.J., Zhou, X.: Estimating the most likely space–time paths, dwell times and path uncertainties from vehicle trajectory data: A time geographic method. *Transportation Research Part C: Emerging Technologies* **66**, 176–194 (2016)
- [52] Treiber, M., Hennecke, A., Helbing, D.: Congested traffic states in empirical observations and microscopic simulations. *Physical review E* **62**(2), 1805 (2000)
- [53] Vemula, A., Muelling, K., Oh, J.: Social attention: Modeling attention in human crowds. In: *2018 IEEE international Conference on Robotics and Automation (ICRA)*. pp. 1–7. IEEE (2018)
- [54] Willard, J., Jia, X., Xu, S., Steinbach, M., Kumar, V.: Integrating physics-based modeling with machine learning: A survey. *arXiv preprint arXiv:2003.04919* (2020)
- [55] Xie, D.F., Fang, Z.Z., Jia, B., He, Z.: A data-driven lane-changing model based on deep learning. *Transportation research part C: emerging technologies* **106**, 41–60 (2019)
- [56] Xu, X., Pang, J., Monterola, C.: Asymmetric optimal-velocity car-following model. *Physica A: Statistical Mechanics and its Applications* **436**, 565–571 (2015)
- [57] Zeng, A., Song, S., Lee, J., Rodriguez, A., Funkhouser, T.: Tossingbot: Learning to throw arbitrary objects with residual physics. *IEEE Transactions on Robotics* (2020)
- [58] Zhan, W., Sun, L., Wang, D., Shi, H., Clausse, A., Naumann, M., Kümmerle, J., Königshof, H., Stiller, C., de La Fortelle, A., Tomizuka, M.: INTERACTION Dataset: An INTERnational, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps. *arXiv:1910.03088 [cs, eess]* (2019)
- [59] Zhang, X., Sun, J., Qi, X., Sun, J.: Simultaneous modeling of car-following and lane-changing behaviors using deep learning. *Transportation research part C: emerging technologies* **104**, 287–304 (2019)
- [60] Ziegler, J., Bender, P., Schreiber, M., Latégahn, H., Strauss, T., Stiller, C., Dang, T., Franke, U., Appenrodt, N., Keller, C.G., et al.: Making bertha drive—an autonomous journey on a historic route. *IEEE Intelligent transportation systems magazine* **6**(2), 8–20 (2014)