



HAL
open science

Interpretable Goal-Based model for Vehicle Trajectory Prediction in Interactive Scenarios

Amina Ghoul, Itheri Yahiaoui, Anne Verroust-Blondet, Fawzi Nashashibi

► **To cite this version:**

Amina Ghoul, Itheri Yahiaoui, Anne Verroust-Blondet, Fawzi Nashashibi. Interpretable Goal-Based model for Vehicle Trajectory Prediction in Interactive Scenarios. IEEE Intelligent Vehicles Symposium (IV), Jun 2023, Anchorage, United States. pp.1-6, 10.1109/IV55152.2023.10186734 . hal-04108657

HAL Id: hal-04108657

<https://hal.science/hal-04108657>

Submitted on 8 Aug 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Interpretable Goal-Based model for Vehicle Trajectory Prediction in Interactive Scenarios

Amina Ghoul¹, Itheri Yahiaoui², Anne Verroust-Blondet¹, and Fawzi Nashashibi¹

Abstract—The abilities to understand the social interaction behaviors between a vehicle and its surroundings while predicting its trajectory in an urban environment are critical for road safety in autonomous driving. Social interactions are hard to explain because of their uncertainty. In recent years, neural network-based methods have been widely used for trajectory prediction and have been shown to outperform hand-crafted methods. However, these methods suffer from their lack of interpretability. In order to overcome this limitation, we combine the interpretability of a discrete choice model with the high accuracy of a neural network-based model for the task of vehicle trajectory prediction in an interactive environment. We implement and evaluate our model using the INTERACTION dataset and demonstrate the effectiveness of our proposed architecture to explain its predictions without compromising the accuracy.

I. INTRODUCTION

Predicting the future motion of a dynamic agent in an interactive environment is crucial many fields and especially in autonomous driving. However, this task is challenging as it depends on various factors such as the agent’s intention or the interaction with his surroundings. Because of these uncertainties, future motion of agents are inherently multimodal. To ensure safe predictions, the agent needs to take into account the dynamics of the surroundings and timely predict their motions in near future to avoid collisions. To address the task of forecasting vehicle motion, many studies use neural network-based model. One major drawback of these methods is the lack of interpretability. In fact, although data-driven approaches achieve outstanding performance in various tasks, it is hard to trust and interpret their predictions. For this reason, in recent years, developing models that can understand social interactions and forecast future trajectories has been an active and challenging area of research.

Early works designed hand-crafted methods based upon domain knowledge to forecast dynamic agents trajectories, either with physics-based models such as Social Forces [1], or with pattern-based models such as discrete choice modelling (DCM) [2]. These models, based on domain knowledge allow their predictions to be interpretable.

The nature of vehicle movement is highly connected to the motion of other road users around them. They alter their paths according to their interactions with neighbors. Thus, the concept of social interaction has been highly evaluated

and discussed in the existed studies [3]. Our interest in this problem stems from the fact that while interaction modeling has been well-investigated in existing studies, it’s hard to interpret the learned social interactions. In these previous studies, variables in models are designed to learn latent behavioral characteristics and with no expectation to have practical implications. For example, the pooling methods [3] directly aggregate hidden states of all neighbors in a neighborhood to learn the connections between people. Thus, it’s hard to understand what kind of social interactions is going on, how it varies among moving pedestrians and how it affects the future trajectories. Attention mechanism [4] can show the interests of agents in each neighbor by observing the learned distribution of attention, thus we know which agent have the greatest influence on agent. However, we still can’t get a concrete pattern of the social interaction. Therefore, these neural network-based models suffer from the lack of interpretability regarding the model’s decision-making process.

To address these limitations, we propose to combine an interpretable discrete choice model with a neural network for the task of vehicle trajectory prediction. Our approach presents a way to easily validate NN models in safety critical applications, by using the interpretable pattern-based rules from the DCM. We conduct extensive experimentations on the real-world INTERACTION dataset and we demonstrate the effectiveness of our method, while at the same time providing a rationale behind high-level decisions, an essential component required for safety-critical applications like autonomous systems. We also conduct a comparative study between two discrete choice models.

II. RELATED WORK

A. Knowledge-based Models

Early works address trajectory prediction problem using of knowledge-based methods. [5] use Kalman filter to predict vehicle future trajectory. Discrete choice modelling (DCM) uses a grid for selecting the next action relative to each individual. DCMs have been used to predict pedestrian’s trajectories [6], and also for many applications in various fields such as facial expression recognition [7]. These knowledge-based methods allow interpretable outputs, but they usually fail to capture the complexity of agent-agent interactions and agent-scene interactions. Therefore, they have low prediction accuracy when predicting trajectories.

This work was carried out in the SAMBA collaborative project, co-funded by BpiFrance in the framework of the Investissement d’Avenir Program.
1: INRIA Paris, France `firstname.lastname@inria.fr`
2: CRESTIC, Université de Reims Champagne-Ardenne, Reims, France
`itheri.yahiaoui@univ-reims.fr`

B. Data-driven Models

In order to solve the low accuracy problem of knowledge-based models, in recent years, many studies tackle the task of motion prediction using neural network models [8], [9]. [3] introduced the social LSTM for pedestrian trajectory prediction. They encode the motion of each agent using an LSTM. Then, they extract the interactions between agents by sharing the hidden states between all the LSTMs corresponding to a set of neighboring pedestrians. MHA JAM [8] applies multi-head attention by considering a joint representation of the static scene and surrounding agents. The authors use each attention head to generate a distinct future trajectory to address multimodality of future trajectories. However, these data-driven methods lack the ability to output predictions that can be explained.

C. Interpretable Trajectory Prediction

To address the lack of interpretability in neural network-based models, recent studies focus on adding expert knowledge to deep learning models for trajectory prediction. Neumeier et al. [10] use an autoencoder where the decoder contains expert knowledge to produce an interpretable latent space in a vehicle trajectory prediction model, in a highway environment. Another way to encourage interpretability in trajectory prediction architectures is through discrete modes. For example, Brewitt et al. [11] propose a Goal Recognition method by Interpretable Trees (GRIT) where the “goal” is defined as many kinds of behavioral intentions, such as “straight-on”, “turn left”, “u-turn”, and “stop”, etc. This aims the goal recognition to be interpretable by humans. Kothari et al. [12] learn a probability distribution over possibilities in an interpretable discrete choice model for the task of pedestrian trajectory prediction.

We use a similar approach for the task of vehicle trajectory prediction in an urban environment. However, unlike [12], we first predict the goal and then the whole trajectory for a prediction horizon greater than 1 second.

To the best of our knowledge, we are the first to use a DCM to help model the behavior of vehicles in their interactions with their surroundings.

In this paper we also consider and compare two types of discrete choice models describing the behavior of vehicles.

III. METHOD

A. Problem definition

The goal is to predict the future trajectories of a target agent T : $\hat{Y}_T = (\hat{x}_T^t, \hat{y}_T^t)$ from time $t = t_{obs} + 1$ to $t = t_f$. We have as input of our model the track history of the target agent and the n neighboring agents in a scene defined as $\mathbf{X} = [X_1, X_2, \dots, X_n]$. Each agent i is represented by a sequence of its states, from time $t = 1$ to $t = t_{obs}$. Each state is composed of a sequence of the agent relative coordinates x_i^t and y_i^t , velocity v_i^t , acceleration a_i^t , heading θ_i^t .

$$X_i^t = (x_i^t, y_i^t, v_i^t, a_i^t, \theta_i^t) \quad (1)$$

The positions of each agent i are expressed in a frame where the origin is the position of the target agent at t_{obs} . The y-axis is oriented toward the target agent’s direction of motion and x-axis points to the direction perpendicular to it.

B. Discrete Choice Model

Discrete choice models or DCMs are hand-crafted models used to explain or predict a choice from a set of alternatives K made by a decision-maker. DCMs are knowledge based models that have a high interpretability. However, despite having interpretable outputs, these models suffer from low prediction accuracy. For that reason, [12] proposed a model combining the high interpretability of the DCMs and the high accuracy of the neural network-based model to predict pedestrian’s trajectories. In this paper, we present an architecture that can model the interactions between vehicles and their surroundings.

We use the Random Utility Maximization (RUM) theory [13] that postulates that the decision-maker aims at maximizing the utility relative to their choice. The utility that an agent i chooses an alternative k , is given as :

$$U_{ik} = \sum_d \beta_d b_{dik} + \epsilon_{ik}, \quad (2)$$

where β are the parameters associated with the explanatory variables b that describe the observed attributes of the choice alternative. We assume that the random terms ϵ_{ik} are independently and identically distributed (i.i.d.) follow an Extreme Value Type I distribution with location parameter zero and the scale parameter 1. In our case, we propose and compare two utility functions u_k for an alternative k . These functions are defined and explained in details in Section IV-C. The alternative k corresponds to the target agent’s goal at timestep t_f , extracted from a radial grid, similar to [12].

C. Neural Networks Model

For a target agent T at time t , X_T^t is embedded using a fully connected layer to a vector e_i^t and encoded using an LSTM encoder,

$$h_i^t = LSTM(h_i^{t-1}, e_i^t; W_{enc}), \quad (3)$$

W_{enc} are the weights to be learned. The weights are shared between all agents in the scene.

Then we build a social tensor similar to [8]. We define the interaction space of a target vehicle T as the area centered on its position at t_{obs} and oriented in its direction of motion. We divide this interaction space into a spatial grid of size (M, N) . The trajectory encoder states of the surrounding agents $h_i^{t_{obs}}$ are placed at their corresponding positions in the 2D spatial grid, giving us a tensor F_s of size (M, N, C_h) , where C_h is the size of the trajectory encoder state.

We use the multi-head attention mechanism [14] to model the social interactions, where the target vehicle $h_T^{t_{obs}}$ is processed by a fully connected layer to give the query and the social tensor is processed by 1×1 convolutional layer to give the keys and the values.




Fig. 1. Architecture of the compared methods for trajectory prediction. The models take as inputs the past trajectories of the agents in the scene (MHA-LSTM), the target coordinates sampled from a radial grid (G-MHA-LSTM), as well as the input of the DCM model (DCM-MHA-LSTM). They output L trajectories. For more details see section III.

We consider K attention heads where K attention heads are specialized to the K potential goals.

For each attention head, we concatenate the output of the multi-head attention module A_k with the target vehicle trajectory encoder state $h_T^{t_{obs}}$ to give a context representation z_k for $k = 1, \dots, K$.

$$z_k = \text{Concat}(h_T^{t_{obs}}, A_k) \quad (4)$$

In order to help the knowledge-based model DCM capture the long term dependencies and the complex interactions, we use the Learning Multinomial Logit (L-MNL) [15] framework.

The goal selection probabilities is defined as :

$$\pi(a_k | \mathbf{X}) = \frac{e^{s_k(\mathbf{X})}}{\sum_{j \in K} e^{s_j(\mathbf{X})}}, \quad (5)$$

where

$$s_k(\mathbf{X}) = u_k(\mathbf{X}) + z_k(\mathbf{X}), \quad (6)$$

where $s_k(\mathbf{X})$ represents the goal function containing the NN encoded terms, $z_k(\mathbf{X})$, as well as utility function $u_k(\mathbf{X})$, following the L-MNL framework.

We consider L attention heads, for each attention head, we concatenate the output of the multi-head attention module A_l with the target vehicle trajectory encoder state $h_T^{t_{obs}}$ to give a context representation c_l for $l = 1, \dots, L$.

$$c_l = \text{Concat}(h_T^{t_{obs}}, A_l) \quad (7)$$

We select the L best scored targets, and we concatenate their embedding to the output of the context representation c_l for $l = 1, \dots, L$.

Finally, the context vector c_l is fed to an LSTM Decoder which generates the predicted parameters of the distributions over the target vehicle's estimated future positions of each possible trajectory for next t_f time steps,

$$\Theta_l^t = \Lambda(\text{LSTM}(h_l^{t-1}, z_l; W_{dec})), \quad (8)$$

where W_{dec} are the weights to be learned, and Λ is a fully connected layer. Similar to [8], we also output the probability P_l associated with each mixture component.

D. Loss function

Our proposed model (DCM-MHA-LSTM) outputs the means and variances $\Theta_l^t = (\mu_l^t, \Sigma_l^t)$ of the Gaussian distributions for each mixture component at each time step. The loss for training the model is composed of a regression loss L_{reg} and two classification losses L_{score} and L_{cls} . L_{reg} is the negative log-likelihood (NLL) similar to the one used in [8] and given by :

$$L_{reg} = -\min_l \sum_{t=t_{obs}+1}^{t_{obs}+t_f} \log(\mathcal{N}(y^t | \mu_l^t; \Sigma_l^t)). \quad (9)$$

L_{score} is a cross entropy loss defined as :

$$L_{score} = -\sum_{l=1}^L \delta_{l*}(l) \log(P_l), \quad (10)$$

where δ is a function equal to 1 if $l = l^*$ and 0 otherwise. L_{cls} is also a cross entropy loss defined as :

$$L_{cls} = -\sum_{k=1}^K \delta_{k*}(k) \log(p_k), \quad (11)$$

where p_k is the probability associated with the potential goal k , δ is a function equal to 1 if $k = k^*$ and 0 otherwise, k^* is the index of the potential goal most closely matching the endpoint of the ground truth trajectory.

Finally, the loss is given by :

$$L = L_{cls} + L_{reg} + L_{score}, \quad (12)$$

IV. EXPERIMENTS

A. Dataset

We evaluate our model on the INTERACTION [16] dataset. The INTERACTION dataset provides a large set of challenging intersection, roundabout, and highway merge scenarios. In total, the data is collected from 11 locations using drones or fixed cameras.

B. Compared Methods

The experiment includes a comparison of different models:

- **I) MHA-LSTM [4]:** This model only takes as inputs the past trajectories of the agents in the scene and outputs L trajectories with their associated probabilities (see the architecture in the red rectangle in Fig. 1). We use $L = 6$ attention heads.
- **II) G-MHA-LSTM [17]:** We add to the previous model a radial grid representation from which we extract potential goals. We predict the goal and then the trajectories conditioned on the predicted goal. (see the architecture in the orange rectangle in Fig. 1).
- **III) DCM-MHA-LSTM :** To predict the goal of the target agent, we combine the DCM and the neural network using the LMNL framework [15]. This model is described in Section III and the architecture is illustrated in the blue rectangle in Fig. 1.
- **IV) ODCM-MHA-LSTM :** This model only uses the DCM to predict the goal of the target agent.

Goal set representations : We also compare different types of radial grids. For the methods II), III) and IV), we compare our results for two types of radial grid : a **dynamic** grid (d) and a **fixed** one (f). Similar to [12], we build the dynamic grid by considering the target agent’s current velocity $v_T^{t_{obs}}$. If $v_T^{t_{obs}} = 0$, we replace it with an arbitrary value equals to 0.5 m.s^{-1} . The fixed grid is built using the value $v = 5.83 \text{ m.s}^{-1}$, which corresponds to the mean of the velocities in the INTERACTION training set.

C. Compared DCMs

We compare two types of DCMs for modelling the behavior of vehicle motion. For our case, the functions modelling vehicle motion phenomenon which we consider for goal selection in this work are:

- 1) *occupancy*: directions containing neighbours in the vicinity are less desirable.
 - 2) *keep direction*: vehicles tend to maintain the same direction of motion.
 - 3) *collision avoidance*: when a neighbour vehicle’s trajectory is head-on towards a potential goal, this goal becomes less desirable due to the chance of a collision.
- **1) DCM 1 :** For the first DCM configuration, we use a utility function defined as:

$$u_k(\mathbf{X}) = \beta_{dir}dir_k + \beta_{occ}occ_k + \beta_{col}col_k \quad (13)$$

Where the functions dir_k , occ_k , and col_k correspond respectively to *keep direction*, *occupancy* and *collision avoidance*. These functions are defined in [2] and [6].

- **2) DCM 2 :** For the second DCM, the utility function is defined as :

$$u_k(\mathbf{X}) = \beta_{dir}dir_k + \beta_{occup}occup_k \quad (14)$$

Where the function dir_k is the same as in (IV-C). For $occup_k$, we use the same mathematical formula as the occupancy function in (IV-C), however, we

don’t consider the position of the neighbors at time t_{obs} . Instead, we consider their predicted position at time $t_{obs} + t_f$ using a Constant velocity model. We assume that before predicting his goal, the target agent first predicts the future positions of his surroundings according to their headings and current velocities, and then avoids the zones that are expected to be crowded. While training this model, we calculate the $occup_k$ function using the ground truth positions of the neighbors.

D. Implementation details

We use $K = 15$ number of potential goals. Similar to [8], our interaction space is 40 m ahead of the target vehicle, 10 m behind and 25 m on each side. We consider the neighbors situated in the interaction space at t_{obs} . We also take into account the neighbors that are susceptible of being in this space from time t_{obs} to t_f . To do so, we predict the trajectories of all of the neighbors in the scene using a Constant Velocity model and if they have a predicted position in the interaction space, we consider them in our model. We argue that this representation allows to consider neighbors that are not situated in the grid at t_{obs} but that can appear in the grid from time $t = t_{obs} + 1$ to $t = t_f$. without having to create a bigger interaction space which can be more computationally expensive. We use $L + K = 6 + 15$ parallel attention operations. We use a batch size of 64 and Adam optimizer. The model is implemented using PyTorch [18].

V. RESULTS

A. Evaluation metrics

Our method for trajectory forecasting is evaluated with the following three error metrics:

- **Minimum Average Displacement Error over k ($minADE_k$) :** The average of pointwise L2 distances between the predicted trajectory and ground truth over the k most likely predictions.
- **Minimum Final Displacement Error over k ($minFDE_k$) :** The final displacement error (FDE) is the L2 distance between the final points of the prediction and ground truth. We take the minimum FDE over the k most likely predictions and average over all agents.
- **Collision II - Groundtruth collision (Col-II) [19]:** This metric calculates the percentage of collision between the primary vehicle’s prediction and the neighbors in the groundtruth future scene.

B. Comparison of Methods

We compare the methods described in Section IV-B.

The results are reported in Table I. DCM¹ and DCM² refers to the first (resp the second) type of DCM described in IV-C. (f) and (d) correspond to respectively, the fixed and the dynamic radial grid representation for the extraction of potential goals. We can see that adding the DCM module decrease the percentage of collisions. We can see that the models using a fixed grid perform slightly better than when




Fig. 2. Qualitative illustration of the ability of our architecture to output high-level interpretable goals. The potential goals are shown in black and the predicted goal is shown in magenta. The ground truth trajectory is in red and the predicted trajectory is in cyan. Current neighbour positions are shown in blue and their past trajectories are shown in green. In the first row, the decision of the model is influenced by the neural-network (NN). In the second row, the decision of the model is strongly influenced by the keep direction map of the DCM. (Un)favourable potential goals are shown in green (red).

using a dynamic one. Thus, we can conclude that adding the information about the velocity doesn't improve the results. For future work, we can try multiple grid configurations, which could potentially improve the results. We can see that using the DCM alone gives worst results as this is due to the fact that the DCM without the NN is not able to predict accurately the goal of the target agent, and therefore, is not able to predict accurate trajectories. We can see that when using the second type of DCM IV-C, the results of ADE/FDE are similar to the ones using the first type of DCM, however, we notice that the percentage of collisions is lower, indicating that in this case, the utility function is more appropriate to avoid collisions.

C. Comparison with the state-of-the-art

We compare our approach with the state-of-the-art using the INTERACTION dataset. Our proposed model does not include any map information. In fact, our aim in this paper, is to study the social interactions between the target agent and his surroundings. Therefore, in Table II, are reported the results where we compare our approach with methods that do not use any map information as well such as DESIRE [20] and Multipath [21]. We can see that our method outperforms these two models. We then compare our approach against state-of-the-art methods that use map information SAN [22], TNT [9], ITRA [23] and ReCoG [24]. The results are reported in Table. III. The main scope is not to compare the approach to the currently best performing trajectory prediction networks. The scope here is to introduce a discrete choice model that provides interpretability, show its feasibility and evaluate the potential of its prediction performance. Nonetheless, our method still achieves competitive results against these methods. Our model is able to perform well while, unlike any of these methods, providing interpretability.

D. Interpretable outputs

1) *Estimation of β* : We study the coefficients β of the utility function of our DCM obtained by training our model.

TABLE I
COMPARISON OF DIFFERENT METHODS ON THE INTERACTION
VALIDATION SET (3 SECS HORIZON)

Model	$MinADE_6$	$MinFDE_6$	$Coll - II$
MHA-LSTM	0.23	0.69	6.1 %
G-MHA-LSTM (f)	0.21	0.58	1.5 %
G-MHA-LSTM (d)	0.22	0.63	1.9 %
ODCM ¹ -MHA-LSTM (f)	0.20	0.58	1.4 %
ODCM ¹ -MHA-LSTM (d)	0.24	0.68	1.0 %
ODCM ² -MHA-LSTM (f)	0.24	0.67	1.2 %
ODCM ² -MHA-LSTM (d)	0.31	0.81	1.3 %
DCM ¹ -MHA-LSTM (f)	0.19	0.57	1.4 %
DCM ¹ -MHA-LSTM (d)	0.20	0.60	1.4 %
DCM ² -MHA-LSTM (f)	0.21	0.57	1.2 %
DCM ² -MHA-LSTM (d)	0.22	0.61	1.1 %

TABLE II
COMPARISON WITH APPROACHES THAT DO NOT USE THE MAP.

	$MinADE_6$	$MinFDE_6$
DESIRE [20]	0.32	0.88
Multipath [21]	0.30	0.99
Ours	0.19	0.58

TABLE III
COMPARISON WITH APPROACHES THAT USE THE MAP.

	$MinADE_6$	$MinFDE_6$
SAN [22]	0.10	0.29
TNT [9]	0.21	0.67
ITRA [23]	0.17	0.49
ReCoG [24]	0.19	0.66
Ours	0.19	0.58

The estimated parameters of both of the utility functions Eq. IV-C and Eq. IV-C are reported in Table. IV for a fixed radial grid. We can see that the all of the coefficients β_{dir} are negative. This means that the utility of an alternative is going to decrease when its angular position is more decentralised with respect to the current direction, respectively. This is coherent as vehicles tend to keep their direction. The β_{occ} parameters have a positive sign, implying that vehicles tend to prefer nearby spatial zones crowded by agents. This result is not coherent as we would expect the contrary. However, this result can be interpreted as in situations where a lot of agents are moving toward the same destination. The β_{col} parameters have a negative sign. This means that vehicles tend to avoid zones where there are potential colliders. We can see that the coefficients β_{occup} and β_{dir} are negatives.

TABLE IV
ESTIMATION OF β

Model	β_{dir}	β_{col}	β_{occ}	β_{occup}
ODCM ¹ -MHA-LSTM (f)	-6.7	-0.7	1.2	-
ODCM ² -MHA-LSTM (f)	-10.4	-	-	-0.3
DCM ¹ -MHA-LSTM (f)	-2.3	-0.3	0.2	-
DCM ² -MHA-LSTM (f)	-2.4	-	-	-0.1

This is coherent as vehicles tend to keep their direction and avoid occupied zones.

2) *Interpretability of the Goals*: We demonstrate the ability of our network to output interpretable goals in Fig. 2. In addition to the predictions map, we illustrate the activation maps of : the neural network (NN) map, the overall DCM map and finally the DCM function for the first DCM configuration described in Section. IV-C. In the first row of Fig. 2, the NN map influence the most the final decision. However, in the second row the influence of the NN map is weaker and the decision is more influenced by the DCM map. We can see that the decision is mostly influenced by keep direction and collision avoidance. Especially, the collision avoidance map helps counteract the influence of the NN map and avoid the model of making the wrong decision. Thus, we observe that the DCM maps work well in conjunction with the NN map to provide interpretable outputs.

VI. CONCLUSION AND FUTURE WORK

In this paper we proposed an interpretable goal based method for the task of vehicle trajectory prediction. In this approach, the discretized goals are selected using both interpretable knowledge-based functions and neural network predictions from the scene. This method allows to combine the high accuracy of the neural networks while being able to understand which vehicle motion rules are present in predicting its goal. Through experiments on the INTERACTION dataset, we highlight the interpretability as well as the accurate predictions outputted by our model. As future work, we plan to add lane information to the radial grid in order to make our model more scene-compliant. Moreover, we plan to explore other types of utility functions for the DCM to better model the behaviour of vehicles in an interactive environment.

REFERENCES

- [1] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.
- [2] G. Antonini, "A discrete choice modeling framework for pedestrian walking behavior with application to human tracking in video sequences," EPFL, Tech. Rep., 2005.
- [3] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 961–971.
- [4] K. Messaoud, I. Yahiaoui, A. Verroust-Blondet, and F. Nashashibi, "Attention based vehicle trajectory prediction," *IEEE Transactions on Intelligent Vehicles*, vol. 6, no. 1, pp. 175–185, 2020.
- [5] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical review E*, vol. 62, no. 2, p. 1805, 2000.
- [6] T. Robin, G. Antonini, M. Bierlaire, and J. Cruz, "Specification, estimation and validation of a pedestrian walking behavior model," *Transportation Research Part B: Methodological*, vol. 43, no. 1, pp. 36–56, 2009.

- [7] G. Antonini, M. Sorci, M. Bierlaire, and J.-P. Thiran, "Discrete choice models for static facial expression recognition," in *Advanced Concepts for Intelligent Vision Systems: 8th International Conference, ACIVS 2006, Antwerp, Belgium, September 18-21, 2006. Proceedings 8*. Springer, 2006, pp. 710–721.
- [8] K. Messaoud, N. Deo, M. M. Trivedi, and F. Nashashibi, "Trajectory prediction for autonomous driving based on multi-head attention with joint agent-map representation," in *IEEE Intelligent Vehicles Symposium, IV 2021, Nagoya, Japan, July 11-17, 2021*. IEEE, 2021, pp. 165–170.
- [9] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid, C. Li, and D. Anguelov, "TNT: target-driven trajectory prediction," in *4th Conference on Robot Learning, CoRL 2020*, ser. Proceedings of Machine Learning Research, vol. 155, 2020, pp. 895–904.
- [10] M. Neumeier, M. Botsch, A. Tollkühn, and T. Berberich, "Variational autoencoder-based vehicle trajectory prediction with an interpretable latent space," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 820–827.
- [11] C. Brewitt, B. Gyevar, S. Garcin, and S. V. Albrecht, "Grit: Fast, interpretable, and verifiable goal recognition with learned decision trees for autonomous driving," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1023–1030.
- [12] P. Kothari, B. Siffringer, and A. Alahi, "Interpretable social anchors for human trajectory forecasting in crowds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 15 556–15 566.
- [13] C. F. Manski, "The structure of random utility models," *Theory and decision*, vol. 8, no. 3, p. 229, 1977.
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [15] B. Siffringer, V. Lurkin, and A. Alahi, "Enhancing discrete choice models with representation learning," *Transportation Research Part B: Methodological*, vol. 140, pp. 236–261, 2020.
- [16] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kummerle, H. Königshof, C. Stiller, A. de La Fortelle *et al.*, "Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps," *arXiv preprint arXiv:1910.03088*, 2019.
- [17] A. Ghoul, K. Messaoud, I. Yahiaoui, A. Verroust-Blondet, and F. Nashashibi, "A lightweight goal-based model for trajectory prediction," in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2022, pp. 4209–4214.
- [18] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.
- [19] P. Kothari, S. Kreiss, and A. Alahi, "Human trajectory forecasting in crowds: A deep learning perspective," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 7386–7400, 2021.
- [20] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker, "Desire: Distant future prediction in dynamic scenes with interacting agents," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 336–345.
- [21] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," *arXiv preprint arXiv:1910.05449*, 2019.
- [22] F. Janjoš, M. Dolgov, M. Kurić, Y. Shen, and J. M. Zöllner, "San: Scene anchor networks for joint action-space prediction," in *2022 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2022, pp. 1751–1756.
- [23] A. Šcibior, V. Lioutas, D. Reda, P. Bateni, and F. Wood, "Imagining the road ahead: Multi-agent trajectory prediction via differentiable simulation," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 720–725.
- [24] X. Mo, Y. Xing, and C. Lv, "Recog: A deep learning framework with heterogeneous graph for interaction-aware trajectory prediction," *arXiv preprint arXiv:2012.05032*, 2020.