

Studying the Effect of Explicit Interaction Representations on Learning Scene-level Distributions of Human Trajectories

Anna Mészáros*, Javier Alonso-Mora, and Jens Kober

Abstract—Effectively capturing the joint distribution of all agents in a scene is relevant for predicting the true evolution of the scene and in turn providing more accurate information to the decision processes of autonomous vehicles. While new models have been developed for this purpose in recent years, it remains unclear how to best represent the joint distributions particularly from the perspective of the interactions between agents. Thus far there is no clear consensus on how best to represent interactions between agents; whether they should be learned implicitly from data by neural networks, or explicitly modeled using the spatial and temporal relations that are more grounded in human decision-making. This paper aims to study various means of describing interactions within the same network structure and their effect on the final learned joint distributions. Our findings show that more often than not, simply allowing a network to establish interactive connections between agents based on data has a detrimental effect on performance. Instead, having well defined interactions (such as which agent of an agent pair passes first at an intersection) can often bring about a clear boost in performance.

I. INTRODUCTION

Autonomous vehicles (AVs) promise to bring about safer and more accessible roads [1], [2]. However, for an AV to navigate in an environment with people, from drivers to vulnerable road users such as cyclists and pedestrians, AVs require the capacity to anticipate what the people around them will do. Predicting human behavior in traffic, however, is complicated by the fact that such behavior is generally not deterministic, but stochastic, taking the form of potentially complex and multi-modal distributions [3].

A possible way to approach predicting the future trajectories of agents is by predicting each individual agent in the scene based on context information such as the past trajectories of agents in the scene and static environment information [4]. These types of prediction models provide what are formally known as marginal predictions. However, a drawback of these types of models is that they fail to account for interactions between agents in the future, resulting in misinterpretations of the actual likelihood of certain outcomes – e.g. two vehicles colliding – when observing the evolution of a scene as a whole [5]. Recently, new prediction models have been developed to address the matter of joint predictions of future trajectories of all agents in a scene. A number of these models leverage multi-head attention or transformer networks to capture the interaction between

agents and jointly predict their future trajectories [6], [7]. Meanwhile, other models utilize learned interaction graphs to not only capture interactions between agents but also to factorize the joint distribution and in turn simplify the learning process such as in the FJMP model [8].

The manner in which the interactions between agents should be provided in machine learning approaches remains under-explored, despite this being one of the common influencing factors used to guide human trajectory prediction models [9]. Very often, these interactions are captured implicitly by the network itself [6], [7]. In other cases, agents are connected to each other through fully connected graphs in which the edges carry interaction information in the form of the agents’ relative distance, velocity, or direction [10], [11], [12]. Other works have aimed to capture interactions by constructing cliques between agents [13], [14]. These approaches, however only jointly predict agents within cliques and disregard potential interactions between cliques, and while certain works aim to address this issue, they have thus far remained constrained to modeling homogeneous groups (typically pedestrians) [15]. Meanwhile other approaches leverage spatial heuristics such as the crossing of future trajectories [8], [16], hypothetical future conflict [17], [18], risk of collision [19], direction of movement (i.e. whether agents are moving towards or away from each other) [20], Euclidean distance [21], [22], or the social forces model [23] to provide more structure for learning these interactions. Additionally, more abstract heuristics such as the similarity [24] or the correlation of features [25] have also been used to guide the learning of interactions between agents. Nevertheless, to our knowledge, no work has yet investigated which interaction representation is the most beneficial for machine learning models for human trajectory prediction, particularly for heterogeneous agents.

In this study we also care about how the chosen interaction representation will affect not only how well the model can approximate the individual ground truth trajectories, but also how well the given model can fit distributions over multiple potential future trajectories for a given situation. A promising group of models for capturing complex and even multi-modal distributions from data are Normalizing Flows (NFs) [26]. This family of methods has already been successfully applied for predicting marginal distributions over agent trajectories [27], [28]. Nevertheless, expanding these approaches for predicting joint distributions of all agents in a scene is not straightforward due to the bijective transformation functions that make up the core of NFs which require the dimensionality of the input to a NF to be kept constant. This

*Corresponding author

This research was supported by NWO-NWA project “Acting under uncertainty” (ACT), NWA.1292.19.298.

All authors are with the Cognitive Robotics Department, TU Delft, 2628 CB Delft, The Netherlands {A.Meszaros, J.AlonsoMora, J.Kober}@tudelft.nl

results in complications since the number of agents in a scene often changes. A naive solution to this problem would be to introduce dummy agents and only focus on a fixed number of agents in a scene. However, this can lead to a number of issues, such as wasted computational resources if too many dummy agents are used to ensure the maximum number of agents is covered. Alternatively, one could use heuristics to determine the most important agents to predict, which may however result in relevant agents being overlooked, particularly in more crowded scenarios. In response to these challenges we leverage the idea of factorizing the joint distribution in accordance to a directed acyclic graph (DAG) which captures the direction of interaction between agents.

The contributions of our work are two-fold. Firstly, we propose a Graph-based Motion Prediction (GMoP) model structure based on normalizing flows to capture joint distributions over the trajectories of all agents in a scene which remains flexible to the varying number of agents from scene to scene. Secondly, we use this model structure to study the manner in which interactions between heterogeneous agents¹ might be modeled on the level of the DAG and how these affect the final learned distributions. We perform the study on four popular real-world driving datasets (Argoverse [29], INTERACTION [30], nuScenes [31], and roundD [32]).

II. BACKGROUND: NORMALIZING FLOWS

Normalizing Flows are a generative method, capable of capturing complex distributions. NFs are based on the concept of transforming distributions through a series of differentiable bijective functions into a simple known “base” distribution Z_0 – most commonly a standard normal distribution. One type of flow model is the conditional autoregressive NF [33], consisting of a series of K normalizing layers, and conditioned by a feature vector C that can take the form of, e.g., an encoding of observations like past trajectories, static environment, and social interactions.

In the normalizing direction, the normalizing layers make up the normalizing flow function F which transforms samples z_K from the desired distribution Z_K into the base distribution Z_0 . In the generative direction, it is then possible to draw a sample $z_K = F^{-1}(z_0)$ from the desired non-normal distribution over outputs Z_K , using a sample z_0 from Z_0 . The Probability Density Function (PDF) p_K of Z_K can then also be obtained in terms of the PDF p_0 :

$$\begin{aligned} p_K(z_K) &= p_0(F(z_K)) |\det J_F(z_K)| \\ &= p_0(z_0) |\det J_{F^{-1}}(z_0)|^{-1}, \end{aligned}$$

The absolute determinant of the Jacobian $|\det J_F(z_K)|$ quantifies the relative change of volume within a small neighborhood of z_K when transforming it to a sample z_0 using F . This ensures that the probability mass remains the same between the two distributions. The parameters of F are learned by minimizing the negative log-likelihood over a finite number N of training samples $z_{K,n}$:

¹scenes can include a combination of up to four different agent types – vehicles, motorcyclists, cyclists and pedestrians

$$\mathcal{L} \approx -\frac{1}{N} \sum_{n=1}^N \log p_0(F(z_{K,n})) + \log |\det J_F(z_{K,n})|. \quad (1)$$

III. GMoP – GRAPH-BASED MOTION PREDICTION

The goal of our approach is that given the past observations of all agents in a scene $\mathbf{X} \in \mathbb{R}^{n_A \times n_I \times 2}$, where n_A is the number of agents and n_I is the number of input timesteps, and information about the static environment in the form of scene graphs - which together form the context information C - to predict a distribution over the future trajectories $\mathbf{Y} \in \mathbb{R}^{n_A \times n_O \times 2}$ of the agents, where n_O is the number of output timesteps.

An important thing to consider when working with normalizing flows is that the number of features need to be preserved between the distribution that one is trying to fit and the base (i.e. latent) distribution of the NF. To circumvent this issue we leverage the idea of factorizing the joint distribution over all agents in a scene based on a learned interaction graph in the form of a directed acyclic graph (DAG) as introduced in FJMP [8]. Factorizing the problem in accordance to a DAG enables us to break the problem of learning a complete joint distribution into learning a series of smaller conditional distributions in the form

$$p_K(\mathbf{Y}|C) = \prod_a^A p_K(\mathbf{y}_a | \mathbf{Y}_{P,a}, C), \quad (2)$$

where $\mathbf{Y}_{P,a}$ are the future trajectories of the parent nodes P of a given agent, a . As a result, the feature dimension of z_0 will always be equal to that of $\mathbf{y}_a$ (where $\mathbf{y}_a = z_K$) for each conditional distribution $p_K(\mathbf{y}_a | \mathbf{Y}_{P,a}, C)$, irrespective of the number of agents in the scene.

To obtain this factorization we first need a DAG, based on an interaction graph for the agents in a given scene. We take this opportunity to investigate the manner in which the interaction graph can be constructed.

A. Learning the Interaction Graph

For predicting the interaction graph we draw inspiration from FJMP [8] and employ a simple two layer MLP classifier ϕ_{interCl} with an embedding layer $\phi_{\text{em}} : \mathbb{R}^f \rightarrow \mathbb{R}^M$ that has an input feature dimensionality of f and embedding dimensionality M , and a classification layer $\phi_{\text{cl}} : \mathbb{R}^M \rightarrow \mathbb{R}^3$ followed by a Softmax activation layer to ensure the probabilities of the three classes (m influences n , n influences m , and no interaction) sums to one.

As input to the classifier we provide the encoded past trajectories of an m - n agent-pair $\hat{\mathbf{x}}_m$, and $\hat{\mathbf{x}}_n$ - obtained by passing the past trajectories $\mathbf{x}_m$ and $\mathbf{x}_n$ through an RNN $\phi_{\text{RNN}_{\text{inter}}}$ - a one-hot encoding of the agents’ types, a distance vector $d_{m,n} = x_{H_m} - x_{H_n}$, and lastly the angle $\alpha_{m,n}$ between the agents’ distance vector $d_{m,n}$ and the source agent’s (agent m in this case) final displacement $\hat{\mathbf{x}}_m = x_{H_m} - x_{H-1_m}$ obtained as

$$\alpha = \text{acos} \left(\frac{d_{m,n} \cdot \hat{\mathbf{x}}_m}{\|d_{m,n}\| \cdot \|\hat{\mathbf{x}}_m\|} \right).$$




Fig. 1. Architecture of GMoP. During training we use the normalizing direction in which we encode the future trajectories $\mathbf{Y}$ with E_{RNN} and transform the abstracted features $\mathbf{y}_a$ to a sample $\mathbf{z}_{0,a} = F(\mathbf{y}_a)$ assumed to follow a standard normal distribution with the probability density function p_0 . For inference we then use the generative direction, in which a sample $\mathbf{z}_{0,a} \sim p_0$ is inversely transformed by the Normalizing Flow to generate the abstracted future trajectories $\mathbf{y}_a = F^{-1}(\mathbf{z}_{0,a})$ that are decoded with D_{RNN} into the actual trajectories $\mathbf{Y}_{\text{pred}}$. The likelihood of the encoded scene-level predictions is obtained with $\prod_a^A p_0(\mathbf{z}_{0,a}) |\det J_{F^{-1}}(\mathbf{z}_{0,a})|^{-1}$. The encoding ϕ_{CNN} of map E , and the scene graph processing modules LaneGCN and Map-to-Agent are optional modules, which can provide richer context information.

The angle α was defined in such a manner to capture interactions between agent pairs in terms of whether agent m is approaching or moving away from agent n . The agent pair information, observed from the perspective of agent m , is concatenated into a single vector $i_{m,n} \in \mathbb{R}^f$. This is then passed through the classifier $cl_{m,n} = \phi_{\text{interCL}}(i_{m,n}) = \phi_{\text{cl}}(\phi_{\text{em}}(i_{m,n}))$ to obtain the classification $cl_{m,n}$ of the interaction between the m - n agent-pair. An overview of the complete architecture used can be found in Fig. 1.

B. Interaction Graph Heuristics

a) **Euclidean distance:** is one of the most commonly used heuristics for establishing interactions between agents. Agents are considered to be interacting if they are within a distance ϵ of each other in the last past timestep. Since we need uni-directional edges between agents for factorizing the joint distribution, we expand this heuristic to consider the angle $\phi_{m,n}$ at which an agent m sees an agent n . This angle is obtained as:

$$\phi_{m,n} = |\text{atan2}(s_{n,y} - s_{m,y}, s_{n,x} - s_{m,x}) - \gamma_m| \in [0, \pi] \quad (3)$$

where $s_{m,x}$, $s_{n,x}$, $s_{m,y}$, $s_{n,y}$ represent the x and y positions of agents m and n at the last past timestep H , and γ_m is the heading angle of agent m . We use an angle of $\phi_{m,n} \in [0, \pi]$ since we do not distinguish if an agent is seen on the left or right from the view of the observing agent m . An agent n is influenced by an agent m , if the angle $\phi_{n,m}$ at which agent n sees agent m is smaller than $\phi_{m,n}$. This then translates to a directed edge going from agent m to n . We do this under the assumption that an agent focuses more on another agent the closer the other agent is to the center of their field of view, i.e. aligned with their heading direction. We further provide the strength of the interaction in the form of edge weights defined as:

$$w_{m,n} = \frac{\epsilon - d(x_{H_m}, x_{H_n})}{\epsilon}, \quad (4)$$

where $d(x_{H_m}, x_{H_n})$ is the Euclidean distance between the positions x_{H_m} and x_{H_n} of agents m and n at the last past timestep H . This weighing captures the aspect that the further away an agent m is from agent n , the less of an effect

this agent has on agent n . It should be noted that for this heuristic, no training is needed, since the construction of the graph follows directly from the above formalization.

b) **The crossing heuristic:** used in FJMP [8], determines the type of interaction between agent pairs from the agents' ground truth future trajectories. More concretely, they perform a crossing check between pairs of trajectories. If agent m reaches the crossing point before agent n , then agent m influences agent n . If there is no crossing point between the trajectories then there is no influence and in turn no edge between the two agents in the interaction graph. This information was then used to pre-train the interaction graph.

In our approach we formalize this heuristic in the following manner. Given the trajectory pairs $(\mathbf{y}_m, \mathbf{y}_n)$ we calculate a distance matrix D of the form:

$$D = \begin{bmatrix} d(y_{0m}, y_{0n}) & \dots & d(y_{0m}, y_{Tn}) \\ \vdots & \ddots & \vdots \\ d(y_{Tm}, y_{0n}) & \dots & d(y_{Tm}, y_{Tn}) \end{bmatrix} \quad (5)$$

for all timesteps of the future trajectories, where $d(y_{tm}, y_{tn})$ is the Euclidean distance between the trajectories $\mathbf{y}_m$ and $\mathbf{y}_n$ at timestep $t \in [0, T]$. To establish whether the trajectories cross at any point in time, we check whether any of the distances in D fall under a specified threshold ϵ_m^a , dependent on the type of agent a and their average width, since the distances are calculated based on agent center-points. More specifically, a crossing is detected when

$$d(y_{tm}, y_{tn}) \leq \epsilon_m^a. \quad (6)$$

We then check the timesteps for the first detected crossing. If $t_{c,m} < t_{c,n}$, then agent m influences agent n which in the interaction graph translates to a directed edge going from agent m to agent n . Conversely if $t_{c,m} > t_{c,n}$, the directed edge points from agent n to agent m . If no crossings are detected, there is no interaction between the agents and as such no edge.

A shortcoming of this heuristic is that it assumes agent paths need to cross for an interaction to occur. However, interactions can also occur in the form of preventing an agent from carrying out their intended behavior. A simple example




Fig. 2. Illustration of the interaction graphs obtained using the (flipped) crossing, (flipped) hypothetical crossing and Euclidean heuristics. The resulting interaction graphs effectively represent factor graphs from which the joint distribution can be obtained. For example, in the case of the hypothetical crossing heuristic the joint distribution is factorized as $P(A)P(B|A)P(C|A)$ while in its flipped version, it is factorized as $P(B)P(C)P(A|B, C)$. Note, for the hypothetical crossing scenario, the hypothetical future is only used to check whether agents' paths would have crossed had there been no interaction between them. The direction of influence, i.e. interaction, is determined by which agent's GT future reached the hypothetical crossing point first.

of this is a pedestrian standing on the side of a road, intending to cross yet being forced to wait due to vehicles passing by.

c) **Hypothetical crossing heuristic:** is a variant of the above heuristic which we define to address the aforementioned shortcoming. For this we extrapolate a given ground truth future trajectory based on the agent's velocity $\dot{x}_H$ and heading θ_H at the last timestep of the past trajectory $\mathbf{x}$. To ensure realistic trajectories, we use the ground truth future trajectory $\mathbf{y}$ as a guideline, so as to maintain the overall shape of the trajectory while speeding it up. The manner in which the speed up is performed is based on two factors. The first is a check for $\dot{x}_H^a < v^a$, where v^a is the average speed for a given agent type a . The average speeds for the agent types are based on the statistics obtained from the nuScenes [31] dataset as a guideline for normal speeds in urban environments. This check captures both agents which are at a standstill as well as agents forced to move at a slower pace than normal. For these agents, we extrapolate the trajectories such that

$$\forall t \dot{y}_t^a < v^a \Rightarrow \dot{y}_t^a = v^a.$$

This ensures, that if at any point in the future the agent's trajectory moves at less than average velocity, it will be sped up to reach the average velocity for its agent type. In the case that $\dot{x}_H^a > v^a$, the extrapolation is performed such that

$$\forall t \dot{y}_t^a < \dot{x}_H^a \Rightarrow \dot{y}_t^a = \dot{x}_H^a.$$

This captures the cases where an agent was forced to slow down in the future, potentially due to surrounding agents. Based on these extrapolated trajectories, a crossing check such as the one performed for the crossing heuristic is applied. However, the result of the crossing check is not directly used to define the interactions. Instead it is only used as an indication of whether the agents' trajectories might have crossed had the two agents ignored each other. If so, we then check which agent reached the crossing point first in reality, and use this to define the direction of interaction analogously to the definition of the crossing heuristic.

d) **Flipped crossing and flipped hypothetical crossing heuristics:** are variants of the crossing and hypothetical crossing heuristics in which we flip the direction of the interaction edges. From a mathematical perspective, both would be valid factorizations of a joint distribution since $P(A, B) = P(A|B)P(B) = P(B|A)P(A)$. What we aim

to study is whether the semantic meaning of the interaction edges can affect the training of the network.

In the case of either the **(flipped) crossing** or **(flipped) hypothetical crossing heuristic** the interaction classifier (Sec. III-A) is pretrained using the interaction classes for every agent-pair according to the used heuristic. A weighted cross-entropy loss is used to account for class imbalance between interacting and non-interacting agent-pairs.

$$\mathcal{L} = - \sum_c^C w_c y_c \log(p_c)$$

where w_c is the class weight, y_c the true class probability which is set to 1 for the correct class, and p_c is the predicted class probability. After training, the interaction classifier and the RNN-encoder $\phi_{\text{RNN}_{\text{inter}}}$ for the past trajectories are frozen.

e) **No heuristic:** is a further case we investigate since there is always the potential that a chosen heuristic is not the correct one for a given task. In this case, the neural network has to learn the best graph connections purely from data in accordance to the loss for fitting the final distribution as provided in Sec. III-C.

Additional information in the form of edge weights for the four crossing heuristics and for the case of no heuristics is provided based on the predicted classification probabilities for the three possible cases (no interaction, agent m affects agent n , and agent n affects agent m). The structure of the interaction classifier is described in Sec. III-A.

f) **Independence between agents:** is an assumption used as a baseline within our method. Although we assume there are no interactions in this case, the training loss from Sec. III-C still optimizes for fitting a joint distribution.

C. Fitting the Distribution

It was previously shown that a better fitting over distributions can be achieved by leveraging an auto-encoder to obtain an abstracted representation of the trajectories over which a distribution should be learned [28]. We apply the same strategy of first transforming the trajectories from positions to displacements with $\tilde{y}_t = y_t - y_{t-1}$, and then encoding these trajectories using the RNN-encoder E_{RNN} of a pre-trained RNN-auto-encoder (RNN-AE) to obtain abstracted trajectories $\hat{\mathbf{y}}$. The abstracted trajectories are then passed to the NF. The conditional distribution being learned is thus $p_K(\hat{\mathbf{y}}_a | \hat{\mathbf{y}}_{P,a}, C)$. For our specific implementation, we

aggregate the future predictions of the parent nodes of agent a through sum-pooling to obtain $\hat{Y}_{P,a}$.

Since we are fitting a joint distribution, we employ a loss over the joint distribution of all agents in a scene:

$$\mathcal{L} = -\frac{1}{N} \sum_{n=1}^N \sum_a^A \log p_0(F(\hat{\mathbf{y}}_{a,n})) + \log |\det J_F(\hat{\mathbf{y}}_{a,n})| \quad (7)$$

During the final inference, the predicted abstracted trajectories $\hat{\mathbf{y}}$ are decoded using the RNN-decoder D_{RNN} of the RNN-AE to obtain the final predicted trajectory $\mathbf{y}_{\text{pred}}$.

D. Encoding Context Information

For encoding the context information we use a series of neural networks. First, we use an RNN encoder ϕ_{RNN} to encode the past trajectories $\mathbf{X}$ of all agents in the scene, giving us $\hat{\mathbf{X}}$. If a scene-graph G_S is provided by the dataset, this graph is processed by LaneGCN [34] and fused with the encoded past trajectories $\hat{\mathbf{X}}$ using the Map-to-Agent network from [34]. The enriched encoding is then further passed through a Graph Neural Network (GNN) ϕ_{GNN} to encode the interactions between the agents in the scene. For this GNN we utilize the same graph structure as the one provided by the previously described interaction heuristics. The output of the GNN is the resulting context vector $\mathbf{C} = \phi_{\text{GNN}}(\hat{\mathbf{X}})$.

IV. EXPERIMENTAL SETUP

A. Models

To study the effect of the interaction graph's structure on the final prediction performance we compare the seven versions of our method:

- GMoP assuming independence between agents (indep)
- GMoP without heuristics (w/o H)
- GMoP with Euclidean distance heuristic (euclH)
- GMoP with crossing heuristic (crH)
- GMoP with hypothetical crossing heuristic (hcrH)
- GMoP with flipped crossing heuristic (flip crH)
- GMoP with flipped hypothetical crossing heuristic (flip hcrH).

We use three state-of-the-art scene-level trajectory prediction models as baselines:

- ADAPT [6] - efficient joint trajectory prediction model.
- AutoBots [7] - transformer based joint prediction model.
- FJMP [8] - an architecture which learns a directed acyclic interaction graph to factorize the future joint distribution for a given scene.

It is important to note that ADAPT and FJMP only provide a fixed number of deterministic joint predictions, where each prediction is a mode. Meanwhile, AutoBots outputs a Mixture Model from which predictions can be sampled.

B. Datasets

We test the methods on four widely used naturalistic datasets with heterogeneous agents, Argoverse [29], INTERACTION [30], nuScenes [31], and roundD [32].

For testing the models on **Argoverse**, we use $n_I = 50$ input timesteps and $n_O = 60$ output timesteps with a

sampling frequency of 10 Hz, resulting in 5 s and 6 s of past and future data respectively. For **INTERACTION**, we use $n_I = 10$ and $n_O = 30$ with a sampling frequency of 10 Hz, resulting in 1 s of past and 3 s of future data. For **nuScenes**, we use $n_I = 4$ and $n_O = 12$ with a sampling frequency of 2 Hz, resulting in 2 s of past and 6 s of future data.

For the three datasets above, the datasets were split based on their respectively provided training and validation sets.

Lastly, for the **roundD** dataset, we use $n_I = 15$ and $n_O = 25$ with a sampling frequency of 5 Hz, resulting in 3 s and 5 s of past and future data respectively. For our experiments, we use the scenarios extracted from the original dataset as done in [35], which focused on the gap acceptance scenario of a vehicle entering the roundabout. This in turn results in a more interactive subset of the original dataset. For this dataset, we employ a leave-one-out training paradigm, in which models would be trained on all but one recording location. This results in three unique splits, one for each of the three locations provided in roundD.

To decrease the effect of random parameter initialization on our final analysis of the results, all of the models are trained using 5 different random seeds.

C. Metrics

To evaluate the predicted trajectories w.r.t. the ground truth we use the following metrics:

- **joint minADE/joint minFDE** - Average/Final L_2 distance (measured in meters) between the best-predicted trajectories for all agents in the scene and the ground truth, based on 6 predicted samples. The number of samples was chosen based on the fact that both ADAPT and FJMP provide 6 predictions. Therefore, for the sake of comparability the lower bound of 6 samples was taken. We chose this metric primarily to obtain an interpretable measure of how closely the predictions of the models capture the single ground truth sample available in real-world test cases, since the usefulness of distribution specific metrics such as Negative Log-Likelihood (NLL) is limited when evaluating on singular ground truth samples. Furthermore, since two of the state-of-the-art methods only predict a fixed number of deterministic trajectories instead of distributions this metric also allows us to better compare our methods to the state-of-the-art.
- **joint NLL** - The average NLL of the ground truth, based on the joint distribution of predicted trajectories for all agents in the scene. This metric provides us with insight on how well the learned distributions fit the ground truth data. To obtain the density estimates needed for the NLL metric, we use a non-parametric density estimation approach proposed in [36] so as to ensure a more reliable comparison between models. For estimating the predicted trajectory distribution, up to 100 sampled trajectories were used.

For the sake of brevity the above metrics will simply be referred to as minADE, minFDE, and NLL.

TABLE I

ARGOVERSE: AVERAGE RESULTS ACROSS FIVE RANDOM SEEDS. FOR EACH METRIC THE BEST PERFORMING MODEL IS UNDERLINED, WHILE BOLD VALUES HIGHLIGHT THE BEST PERFORMING VERSION OF GMoP.

Models	minADE	minFDE	NLL
ADAPT	<u>1.89</u> ±0.07	<u>4.56</u> ±0.13	4.88e ³ ±357.14
AutoBots	4.84±0.52	6.22±0.45	2.98e ⁵ ±7.04e ⁴
FJMP	2.38±0.77	5.96±2.17	8.08e ⁴ ±6.50e ⁴
GMoP w/o H	4.48±4.56	10.65±10.24	146.69±324.02
GMoP indep	2.21±0.05	5.53±0.11	-142.95±73.12
GMoP euclH	5.74±7.06	13.81±16.64	493.68±979.89
GMoP crH	2.21±0.10	5.48±0.17	-149.62±213.09
GMoP hcrH	2.18 ±0.02	5.44 ±0.07	178.27±358.92
GMoP flip crH	2.23±0.02	5.58±0.05	-161.22 ±136.80
GMoP flip hcrH	2.38±0.05	5.60±0.11	236.73±312.31

TABLE II

INTERACTION: AVERAGE RESULTS ACROSS FIVE RANDOM SEEDS. FOR EACH METRIC THE BEST PERFORMING MODEL IS UNDERLINED, WHILE BOLD VALUES HIGHLIGHT THE BEST PERFORMING VERSION OF GMoP.

Models	minADE	minFDE	NLL
ADAPT	1.21±0.10	2.52±0.16	5.27e ³ ±204.91
AutoBots	<u>0.55</u> ±0.04	<u>1.49</u> ±0.05	1.83e ³ ±311.50
FJMP	0.66±0.25	1.82±0.61	2.22e ³ ±1.06e ³
GMoP w/o H	0.74±0.12	2.23±0.35	-632.79±9.42
GMoP indep	0.66 ±0.06	2.00 ±0.16	-640.51±4.77
GMoP euclH	0.70±0.10	2.08±0.27	-624.13±2.72
GMoP crH	0.74±0.11	2.28±0.38	-636.38±6.05
GMoP hcrH	0.68±0.05	2.06±0.16	-640.79 ±0.97
GMoP flip crH	0.78±0.13	2.39±0.41	-633.97±5.80
GMoP flip hcrH	0.70±0.06	2.09±0.15	-639.98±1.98

V. RESULTS

Comparing the performance of the different GMoP variants to the state-of-the-art models, we can see that in terms of the distance metrics minADE/minFDE our models fall in the middle of the performance range attained by the three baselines. It is important to note that our models are optimized for achieving the best possible distribution fits across the dataset rather than trying to predict a discrete set of trajectories for a given scenario. This is also reflected in the NLL metric, where the GMoP variants are able to consistently achieve better results, outperforming even the AutoBots model which also optimizes for distribution fit.

Looking more closely at versions of GMoP, we can quickly see that out of the different versions there is no one choice that reliably works well for all situations. Out of the two large-scale datasets - Argoverse and INTERACTION - we can see that on Argoverse (Tab. I), all of the heuristics save for the **Euclidean heuristic** achieve comparable performance in terms of the distance metrics. We however observe a difference in performance when looking at the distribution fit, in which case the **flipped crossing heuristic** achieved the best average performance followed closely by the **crossing heuristic** and the **independence** assumption.

Meanwhile, on the INTERACTION dataset we observe that the variants of our model perform comparably to each other. The worst performance out of the variants, both in terms of average performance as well as higher standard

deviation can be seen for the cases with **no heuristic** as well as the **(flipped) crossing heuristic**. At the same time, some of the best performance – both in terms of average performance as well as low variability across seeds – is achieved by the **(flipped) hypothetical crossing heuristic** which accounts for agents not being able to execute their intentions due to the actions of the agents around them. This is understandable since INTERACTION is a dataset featuring highly interactive and crowded traffic scenarios, meaning that interactions extend beyond the aspect of which agent crosses paths with another agent first. At the same time, it is noteworthy that similar performance is achieved using the **independence** assumption.

Looking now at roundD Location 0, we can see that the heuristics (**Euclidean distance heuristic**, **(flipped) crossing heuristic**, and **(flipped) hypothetical crossing heuristic**) all perform reasonably well in terms of distance metrics. The **(flipped) hypothetical crossing heuristic** additionally achieves better NLL compared to the other heuristics. Meanwhile, the **independence** assumption, while obtaining worse performance in terms of distance metrics, obtained substantially better performance in terms of NLL. Similar observations had been made in [28], where better performance in distance metrics did not always imply a better distribution fit and vice versa. It is worth noting that the generally poorer performance of both our models and the state-of-the-art baselines on Location 0, compared to the other two locations can be attributed to the difference in dataset sizes generated by the splits, with Location 0 having only 971 training instances, while Locations 1 and 2 were trained with 11220 and 11805 instances respectively. Nevertheless, the better performance of the **(flipped) hypothetical crossing heuristic** compared to other versions of our model on this Location, speaks for the fact that adequate heuristics can guide the learning of a model better in the presence of limited training data. At the same time the better distribution fit achieved when using the **independence** assumption indicates that attempting to establish structure through heuristics can be highly detrimental to the model if the heuristics fail to capture the interactions between agents. Meanwhile, on Location 1 we observe comparable performance between variants of our model with the exception of the model with **no heuristics**. This model achieved the worst performance in terms of distance metrics both in regards to the average metric values as well as a much higher level of variability. On Location 2 we see comparable performance between all variants of our model. The main difference in performance on this location is the performance in regards to the distribution fit, with the best NLL values being obtained through the **independence** assumption.

Meanwhile on nuScenes we once more observe that leveraging heuristics for defining the interactions between agents can help improve performance. Specifically, we observe the best performance in terms of distance metrics when using the **(flipped) crossing heuristic** and **(flipped) hypothetical crossing heuristics**, and to a lesser extent when using the **independence** assumption. At the same time, the **(flipped)**

TABLE III

ROUND: AVERAGE RESULTS ACROSS FIVE RANDOM SEEDS FOR THE THREE ROUND LOCATIONS. FOR EACH METRIC THE BEST PERFORMING MODEL IS UNDERLINED, WHILE BOLD VALUES HIGHLIGHT THE BEST PERFORMING VERSION OF GMoP.

Models	Location 0			Location 1			Location 2		
	minADE	minFDE	NLL	minADE	minFDE	NLL	minADE	minFDE	NLL
ADAPT	9.76 $\pm$ 0.38	16.15 $\pm$ 0.60	5.45e ³ $\pm$ 3.10e ³	7.23 $\pm$ 0.46	12.69 $\pm$ 0.95	2.74e ³ $\pm$ 1.02e ³	6.97 $\pm$ 0.27	12.77 $\pm$ 0.86	7.33e ³ $\pm$ 1.00e ⁴
AutoBots	5.38 $\pm$ 0.50	13.30 $\pm$ 1.24	1.71e ⁴ $\pm$ 4.76e ³	3.14 $\pm$ 0.01	8.51 $\pm$ 0.09	8.32e ³ $\pm$ 1.08e ³	<u>2.95</u> $\pm$ 0.18	<u>7.73</u> $\pm$ 0.50	6.54e ³ $\pm$ 1.05e ³
FJMP	<u>4.50</u> $\pm$ 0.49	<u>11.55</u> $\pm$ 1.66	6.46e ³ $\pm$ 1.97e ³	<u>2.74</u> $\pm$ 0.07	<u>8.12</u> $\pm$ 0.24	7.75e ³ $\pm$ 571.68	3.22 $\pm$ 0.37	9.26 $\pm$ 0.71	4.81e ³ $\pm$ 614.46
GMoP w/o H	7.23 $\pm$ 2.54	16.36 $\pm$ 4.27	212.01 $\pm$ 399.98	4.30 $\pm$ 1.96	11.49 $\pm$ 5.54	-65.97 $\pm$ 24.96	3.44 $\pm$ 0.34	9.19 $\pm$ 0.98	-53.17 $\pm$ 15.80
GMoP indep	6.33 $\pm$ 0.43	14.89 $\pm$ 0.96	-52.83 $\pm$ 15.31	3.59 $\pm$ 0.23	9.52 $\pm$ 0.51	-70.32 $\pm$ 15.74	3.37 $\pm$ 0.13	8.91 $\pm$ 0.55	-89.47 $\pm$ 9.36
GMoP euclH	5.91 $\pm$ 0.23	13.84 $\pm$ 0.55	67.07 $\pm$ 64.59	3.30 $\pm$ 0.23	8.68 $\pm$ 0.63	-69.60 $\pm$ 29.09	3.41 $\pm$ 0.11	8.91 $\pm$ 0.33	-62.41 $\pm$ 21.73
GMoP crH	5.89 $\pm$ 0.46	13.78 $\pm$ 0.97	82.36 $\pm$ 144.12	3.48 $\pm$ 0.15	9.14 $\pm$ 0.46	-61.18 $\pm$ 21.68	3.55 $\pm$ 0.21	9.27 $\pm$ 0.38	-43.32 $\pm$ 29.72
GMoP hcrH	5.67 $\pm$ 0.30	13.44 $\pm$ 0.80	-3.79 $\pm$ 12.16	3.54 $\pm$ 0.24	9.29 $\pm$ 0.56	-50.08 $\pm$ 23.40	3.56 $\pm$ 0.13	9.59 $\pm$ 0.43	-51.14 $\pm$ 13.38
GMoP flip crH	6.47 $\pm$ 0.47	15.00 $\pm$ 0.83	126.99 $\pm$ 148.71	3.48 $\pm$ 0.29	9.12 $\pm$ 0.72	-73.36 $\pm$ 17.32	3.55 $\pm$ 0.28	9.22 $\pm$ 0.62	-67.26 $\pm$ 17.48
GMoP flip hcrH	6.16 $\pm$ 0.27	14.69 $\pm$ 0.66	-19.43 $\pm$ 27.93	3.60 $\pm$ 0.21	9.50 $\pm$ 0.51	-55.40 $\pm$ 22.48	3.40 $\pm$ 0.11	8.96 $\pm$ 0.30	-65.80 $\pm$ 8.81

TABLE IV

NUScenes: AVERAGE RESULTS ACROSS FIVE RANDOM SEEDS. FOR EACH METRIC THE BEST PERFORMING MODEL IS UNDERLINED, WHILE BOLD VALUES HIGHLIGHT THE BEST PERFORMING VERSION OF GMoP.

Models	minADE	minFDE	NLL
ADAPT	3.54 $\pm$ 0.21	7.26 $\pm$ 0.15	508.04 $\pm$ 115.78
AutoBots	<u>3.43</u> $\pm$ 0.07	<u>7.06</u> $\pm$ 0.11	1.01e ³ $\pm$ 124.06
FJMP	5.68 $\pm$ 4.97	11.53 $\pm$ 9.61	1.56e ⁵ $\pm$ 3.01e ⁵
GMoP w/o H	5.48 $\pm$ 2.16	12.15 $\pm$ 4.29	446.70 $\pm$ 202.09
GMoP indep	4.43 $\pm$ 0.11	10.19 $\pm$ 0.27	542.31 $\pm$ 264.23
GMoP euclH	5.79 $\pm$ 2.92	13.28 $\pm$ 6.70	371.91 $\pm$ 120.77
GMoP crH	4.31 $\pm$ 0.10	9.80 $\pm$ 0.25	348.83 $\pm$ 81.07
GMoP hcrH	4.28 $\pm$ 0.09	9.74 $\pm$ 0.18	451.60 $\pm$ 111.90
GMoP flip crH	4.28 $\pm$ 0.12	9.77 $\pm$ 0.26	353.31 $\pm$ 63.52
GMoP flip hcrH	4.31 $\pm$ 0.10	9.85 $\pm$ 0.20	510.90 $\pm$ 112.76

crossing heuristic achieve the best distribution fits both in terms of average performance as well as the variability across random seeds, indicating that this heuristic is able to more reliably capture the interactions between agents compared to the other architecture variants.

VI. CONCLUSION & DISCUSSION

Our findings indicate that relying purely on the network to extract interactions between agents based on the training data is often not sufficient for achieving good prediction performance. In fact, GMoP with no heuristics was generally the poorest performing architecture across our different test cases. At the same time, using the independence assumption and relying on the joint NLL loss to guide the learning of the network often lead to some of the better performing models. This is not to say, however, that heuristics are not beneficial for interpreting the interactions between agents. With the correct heuristics, a clear boost in performance can be achieved as shown on the NuScenes dataset, where the (flipped) crossing heuristic resulted in the best performance between the GMoP variants. What is further worth noting is that in the case of little training data, choosing a good heuristic which captures the underlying mechanism of interaction in the given scenario can aid in the learning of the model. We observed an example of this on Round Location 0 when using the (flipped) hypothetical crossing heuristics as opposed to the (flipped) crossing heuristics. What is interesting, is that although the choice of heuristic

can have a strong impact on the final performance of the network, the direction of the established edges in the graph does not greatly impact performance. We can see this in the comparable performance achieved by the (hypothetical) crossing heuristics compared to their flipped counterparts. This indicates that the presence of the edges is more important than the semantic interpretation of their direction. However, none of the heuristics provided the best performance across the different datasets, indicating that these heuristics are highly situation dependent. While the (flipped) crossing heuristics generally resulted in better performing models, it was in some cases outperformed by the (flipped) hypothetical crossing heuristics and independence assumption, such as on INTERACTION and Round Location 0.

This study acts as a first step towards understanding how to best represent interactions between agents and leverage this for joint predictions. Nevertheless, there are aspects which our study does not investigate and would be important directions for future research. For one, we do not consider potential temporal evolution of the interactions themselves. While the (hypothetical) crossing heuristics strive to capture knowledge about the future evolution of the interaction between two agents it remains a simplification of the continuous negotiations and re-assessments of a situation that traffic participants go through with each other. Secondly, we make the assumption that the interactions between different agent types are governed by the same mechanisms. However, this need not be the case particularly when considering interactions between and within vulnerable and non-vulnerable road users. It would therefore be worthwhile to investigate whether there is a unified mechanism underlying all interactions, or whether interactions should be described differently depending on agents' types. Thirdly, we only look at the trajectories of the agents to determine their influence on each other and disregard potential confounders such as traffic rules which can have an important impact on the behavior of the agents. Lastly, while the type of heuristics we used in our study are common to the field of trajectory prediction they are fairly simplified representations of agent interactions. Meanwhile, the field of human behavior prediction has been striving towards modeling the underlying mechanisms behind the interaction between agents [37], [38].

These models, while grounded by human behavior studies, are often geared towards specific scenarios such as car following, intersections, merging, etc. and are often also not suitable for real-time multi-modal predictions which are relevant for motion planning. For this reason, another interesting direction for future research would be to investigate how to combine the benefits of scalability and real-time inference of neural networks with the deeper understanding of human interactions from the field of human behavior modeling.

REFERENCES

- [1] J. S. Brar and B. Caulfield, "Impact of autonomous vehicles on pedestrians' safety," in *IEEE 20th Int. Conf. on Intell. Transp. Syst. (ITSC)*, 2017.
- [2] J. Meyer, H. Becker, P. M. Bösch, and K. W. Axhausen, "Autonomous vehicles: The next jump in accessibilities?," *Research Transp. Econ.*, vol. 62, pp. 80–91, 2017.
- [3] G. M. Ferro and D. Sornette, "Stochastic representation decision theory: How probabilities and values are entangled dual characteristics in cognitive processes," *Plos one*, vol. 15, no. 12, p. e0243661, 2020.
- [4] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Eur. Conf. on Comput. Vis.*, 2020.
- [5] W. Luo, C. Park, A. Cornman, B. Sapp, and D. Anguelov, "Jfip: Joint future prediction with interactive multi-agent modeling for autonomous driving," in *Conf. on Robot Learn.*, pp. 1457–1467, PMLR, 2023.
- [6] G. Aydemir, A. K. Akan, and F. Güney, "Adapt: Efficient multi-agent trajectory prediction with adaptation," in *Proc. IEEE/CVF Int. Conf. on Comput. Vis.*, pp. 8295–8305, 2023.
- [7] R. Girgis, F. Golemo, F. Codevilla, M. Weiss, J. A. D'Souza, S. E. Kahou, F. Heide, and C. Pal, "Latent variable sequential set transformers for joint multi-agent motion prediction," *arXiv preprint arXiv:2104.00563*, 2021.
- [8] L. Rowe, M. Ethier, E.-H. Dykhne, and K. Czarnecki, "Fjmp: Factorized joint multi-agent motion prediction over learned directed acyclic interaction graphs," in *Proc. IEEE/CVF Conf. on Comput. Vis. Pattern Recognit.*, pp. 13745–13755, 2023.
- [9] D. E. Benrachou, S. Glaser, M. Elhenawy, and A. Rakotonirainy, "Use of social interaction and intention to improve motion prediction within automated vehicle framework: A review," *IEEE Trans. on Intell. Transp. Syst.*, vol. 23, no. 12, pp. 22807–22837, 2022.
- [10] L. Li, J. Su, L. Qiu, J. Lian, and G. Guo, "Efin-mp: Explicit future interaction network for motion prediction," *IEEE Trans. on Intell. Veh.*, 2024.
- [11] X. Chen, F. Luo, F. Zhao, and Q. Ye, "Goal-guided and interaction-aware state refinement graph attention network for multi-agent trajectory prediction," *IEEE Robotics Autom. Lett.*, vol. 9, no. 1, pp. 57–64, 2023.
- [12] X. Mo, Z. Huang, Y. Xing, and C. Lv, "Multi-agent trajectory prediction with heterogeneous edge-enhanced graph attention network," *IEEE Trans. on Intell. Transp. Syst.*, vol. 23, no. 7, pp. 9554–9567, 2022.
- [13] Y. Chen, B. Ivanovic, and M. Pavone, "Scept: Scene-consistent, policy-based trajectory predictions for planning," in *Proc. IEEE/CVF conference on computer vision pattern recognition*, pp. 17103–17112, 2022.
- [14] V. K. Nivash and A. H. Qureshi, "Simmf: Semantics-aware interactive multiagent motion forecasting for autonomous vehicle driving," in *2024 IEEE Int. Conf. on Robotics Autom. (ICRA)*, pp. 6813–6819, IEEE, 2024.
- [15] Y. Chen, C. Liu, X. Mei, B. E. Shi, and M. Liu, "Hgcg-gjs: Hierarchical graph convolutional network with groupwise joint sampling for trajectory prediction," in *2022 IEEE/RSJ Int. Conf. on Intell. Robots Syst. (IROS)*, pp. 13400–13405, IEEE, 2022.
- [16] Q. Sun, X. Huang, J. Gu, B. C. Williams, and H. Zhao, "M2i: From factored marginal trajectory prediction to interactive prediction," in *Proc. IEEE/CVF Conf. on Comput. Vis. Pattern Recognit.*, pp. 6543–6552, 2022.
- [17] T. Zhang, M. Fu, Y. Yang, W. Song, and T. Liu, "Edge-enriched graph transformer for multi-agent trajectory prediction with relative positional semantics," *IEEE Trans. on Instrum. Meas.*, 2024.
- [18] P. Bhattacharyya, C. Huang, and K. Czarnecki, "Ssl-interactions: Pretext tasks for interactive trajectory prediction," in *2024 IEEE Intell. Veh. Symp. (IV)*, pp. 1450–1457, IEEE, 2024.
- [19] Q. Xue, F. Yang, S. Li, X. Li, G. Li, and W. Zhang, "Rethink reynolds' rules: flock-inspired network for vehicle trajectory prediction," *The J. Supercomput.*, vol. 81, no. 6, p. 802, 2025.
- [20] S. Wen, H. Wang, D. Liu, Q. Zhangli, and D. Metaxas, "Second-order graph odes for multi-agent trajectory forecasting," in *Proc. IEEE/CVF Winter Conf. on Appl. Comput. Vis.*, pp. 5101–5110, 2024.
- [21] M. Kang, S. Wang, S. Zhou, K. Ye, J. Jiang, and N. Zheng, "Ffinet: Future feedback interaction network for motion forecasting," *IEEE Trans. on Intell. Transp. Syst.*, 2024.
- [22] Y. Yuan, X. Weng, Y. Ou, and K. M. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *Proc. IEEE/CVF international conference on computer vision*, pp. 9813–9823, 2021.
- [23] Z. Mo, Y. Fu, and X. Di, "Pi-neugode: Physics-informed graph neural ordinary differential equations for spatiotemporal trajectory prediction," in *Proc. 23rd Int. Conf. on Auton. Agents Multiagent Syst.*, pp. 1418–1426, 2024.
- [24] D. Zhu, G. Zhai, Y. Di, F. Manhardt, H. Berkemeyer, T. Tran, N. Navab, F. Tombari, and B. Busam, "Ippc-tp: Utilizing incremental pearson correlation coefficient for joint multi-agent trajectory prediction," in *Proc. IEEE/CVF Conf. on Comput. Vis. Pattern Recognit.*, pp. 5507–5516, 2023.
- [25] C. Xu, M. Li, Z. Ni, Y. Zhang, and S. Chen, "Groupnet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning," in *Proc. IEEE/CVF Conf. on Comput. Vis. Pattern Recognit.*, pp. 6498–6507, 2022.
- [26] E. G. Tabak and C. V. Turner, "A family of nonparametric density estimation algorithms," *Commun. on Pure Appl. Math.*, vol. 66, no. 2, pp. 145–164, 2013.
- [27] C. Schöller and A. Knoll, "FloMo: Tractable motion prediction with normalizing flows," in *Int. Conf. Intell. Robot. Syst.*, 2021.
- [28] A. Mészáros, J. F. Schumann, J. Alonso-Mora, A. Zgonnikov, and J. Kober, "Trajflow: Learning distributions over trajectories for human behavior prediction," in *2024 IEEE Intell. Veh. Symp. (IV)*, pp. 184–191, IEEE, 2024.
- [29] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, et al., "Argoverse 2: Next generation datasets for self-driving perception and forecasting," *arXiv preprint arXiv:2301.00493*, 2023.
- [30] 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.
- [31] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. on Comput. Vis. Pattern Recognit.*, 2020.
- [32] R. Krajewski, T. Moers, J. Bock, L. Vater, and L. Eckstein, "The round dataset: A drone dataset of road user trajectories at roundabouts in Germany," in *IEEE 23rd Int. Conf. on Intell. Transp. Syst.*, 2020.
- [33] Y. Lu and B. Huang, "Structured output learning with conditional generative flows," in *Proc. AAAI Conf. on Artif. Intell.*, 2020.
- [34] M. Liang, B. Yang, R. Hu, Y. Chen, R. Liao, S. Feng, and R. Urtasun, "Learning lane graph representations for motion forecasting," in *Comput. Vision-ECCV 2020: 16th Eur. Conf. Glasgow, UK, August 23–28, 2020. Proceedings, Part II 16*, pp. 541–556, Springer, 2020.
- [35] J. F. Schumann, A. R. Srinivasan, J. Kober, G. Markkula, and A. Zgonnikov, "Using models based on cognitive theory to predict human behavior in traffic: A case study," in *2023 IEEE 20th Int. Conf. on Intell. Transp. Syst. (ITSC)*, 2023.
- [36] A. Mészáros, J. F. Schumann, J. Alonso-Mora, A. Zgonnikov, and J. Kober, "Rome: Robust multi-modal density estimator," *Proc. 33rd Int. Jt. Conf. on Artif. Intell.*, 2024, in press.
- [37] S. Kolekar, J. De Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model," *Nat. communications*, vol. 11, no. 1, pp. 1–13, 2020.
- [38] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Modelling communication-enabled traffic interactions," *Royal Soc. open science*, vol. 10, no. 5, p. 230537, 2023.