

# Diverse and Admissible Trajectory Forecasting through Multimodal Context Understanding

Seong Hyeon Park<sup>1</sup>, Gyubok Lee<sup>2</sup>, Jimin Seo<sup>3\*</sup>, Manoj Bhat<sup>4\*</sup>, Minseok Kang<sup>5</sup>,  
Jonathan Francis<sup>4,6</sup>, Ashwin Jadhav<sup>4</sup>, Paul Pu Liang<sup>4</sup>, and  
Louis-Philippe Morency<sup>4</sup>

<sup>1</sup> Hanyang University, Seoul, Korea  
shpark@spa.hanyang.ac.kr

<sup>2</sup> Yonsei University, Seoul, Korea  
glee48@yonsei.ac.kr

<sup>3</sup> Korea University, Seoul, Korea  
jmseo0607@korea.ac.kr

<sup>4</sup> Carnegie Mellon University, Pittsburgh, PA, USA  
{mbhat, jmf1, arjadhav, pli, morency}@cs.cmu.edu

<sup>5</sup> Sogang University, Seoul, Korea  
ahstarwab@sogang.ac.kr

<sup>6</sup> Bosch Research Pittsburgh, Pittsburgh, PA, USA

**Abstract.** Multi-agent trajectory forecasting in autonomous driving requires an agent to accurately anticipate the behaviors of the surrounding vehicles and pedestrians, for safe and reliable decision-making. Due to partial observability in these dynamical scenes, directly obtaining the posterior distribution over future agent trajectories remains a challenging problem. In realistic embodied environments, each agent’s future trajectories should be both *diverse* since multiple plausible sequences of actions can be used to reach its intended goals, and *admissible* since they must obey physical constraints and stay in drivable areas. In this paper, we propose a model that synthesizes multiple input signals from the multimodal world—the environment’s scene context and interactions between multiple surrounding agents—to best model all diverse and admissible trajectories. We compare our model with strong baselines and ablations across two public datasets and show a significant performance improvement over previous state-of-the-art methods. Lastly, we offer new metrics incorporating admissibility criteria to further study and evaluate the diversity of predictions. Codes are at: <https://github.com/kami93/CMU-DATF>.

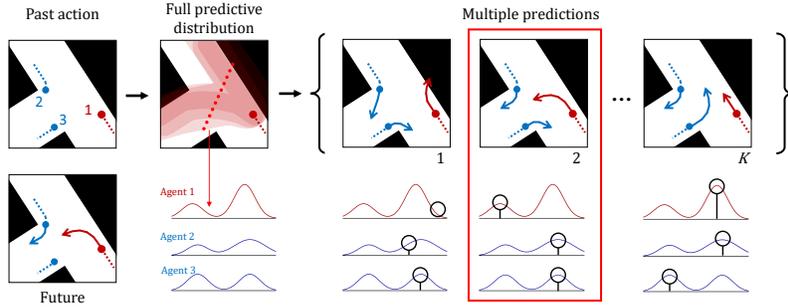
**Keywords:** Trajectory Forecasting · Diversity · Admissibility · Generative Modeling · Autonomous Driving

## 1 Introduction

Trajectory forecasting is an important problem in autonomous driving scenarios, where an autonomous vehicle must anticipate the behavior of other surrounding agents (e.g., vehicles and pedestrians), within a dynamically-changing environment, in order to plan its own actions accordingly. However, since the contexts

---

\* Authors contributed equally

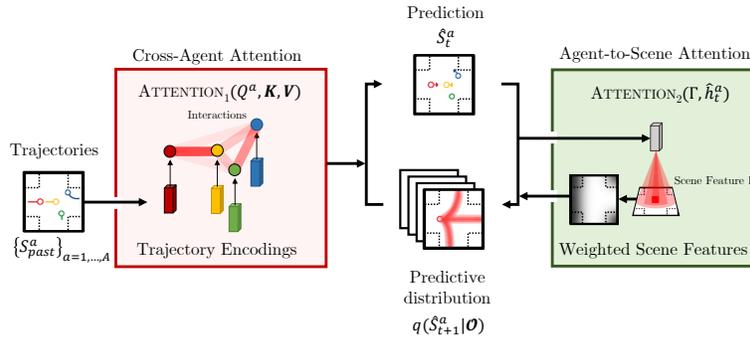


**Fig. 1.** Diverse and admissible trajectory forecasting. Based on the existing context, there can be multiple valid hypothetical future trajectories. Therefore, the predictive distribution of the trajectories should reflect various modes, representing different plausible goals (*diversity*) while penalizing implausible trajectories that either conflict with the other agents or are outside valid drivable areas (*admissibility*).

of agents’ behavior such as intentions, social interactions, or environmental constraints are not directly observed, predicting future trajectories is a challenging problem [22, 26, 33]. It requires an estimation of most likely agent actions based on key observable environmental features (e.g., road structures, agent interactions) as well as the simulation of agents’ hypothetical future trajectories toward their intended goals. In realistic embodied environments, there are multiple plausible sequences of actions that an agent can take to reach its intended goals. However, each trajectory must obey physical constraints (e.g., Newton’s laws) and stay in the statistically plausible locations in the environment (i.e., the drivable areas). In this paper, we refer to these attributes as *diverse* and *admissible* trajectories, and illustrate some examples in Fig. 1. Achieving diverse and admissible trajectory forecasting for autonomous driving allows each agent to make the best predictions, by taking into account all valid actions that other agents could take.

To predict a diverse set of admissible trajectories, each agent must understand its *multimodal* environment, consisting of the scene context as well as interactions between other surrounding agents. Scene context refers to the typical and spatial activity of surrounding objects, presence of traversable area, etc. which contribute to forecasting next maneuvers. These help to understand some semantic constraints on the agent’s motion (e.g., traffic laws, road etiquette) and can be inferred from the present and corresponding multimodal data i.e spatial as well as social, temporal motion behavior data. Therefore, the model’s ability to extract and meaningfully represent multimodal cues is crucial.

Concurrently, another challenging aspect of trajectory forecasting lies in encouraging models to make diverse predictions about future trajectories. Due to high-costs in data collection, diversity is not explicitly present in most public datasets, but only one annotated future trajectories. [7, 16, 18]. Vanilla predictive models that fit future trajectories based only on the existing annotations would severely underestimate the diversity of all possible trajectories. In addition, measuring the quality of predictions using existing annotation-based measures (e.g., displacement errors [23]) does not faithfully score diverse and admissible trajectory predictions.



**Fig. 2.** Overview of our multimodal attention approach. Best viewed in color. The cross-agent attention module (left) generates an attention map over the encoded trajectories of neighboring agents. The agent-to-scene attention model (right) generates an attention map over the scene, based on the drivable-area map.

As a step towards multimodal understanding for diverse and admissible trajectory forecasting, our contributions are *three-fold*:

1. We propose a model that addresses the lack of diversity and admissibility for trajectory forecasting through the understanding of the multimodal environmental context. As illustrated in Fig. 2, our approach explicitly models agent-to-agent and agent-to-scene interactions through “self-attention” [27] among multiple agent trajectory encodings, and a conditional trajectory-aware “visual attention” [31] over the map, respectively. Together with a constrained flow-based decoding, trained with symmetric cross-entropy [21], this allows our model to generate diverse and admissible trajectory candidates by fully integrating all environmental contexts.
2. We propose a new annotation-free approach to estimate the true trajectory distribution based on the drivable-area map. This approximation is utilized for evaluating hypothetical trajectories generated by our model during the learning process. Previous methods [21] depend on ground-truth (GT) recordings to model the real distribution; for most of the time, only one annotation is available per agent. Our approximation method does not rely on GT samples and empirically facilitates greater diversity in the predicted trajectories while ensuring admissibility.
3. We propose a new metric, Drivable Area Occupancy (DAO), to evaluate the diversity of the trajectory predictions while ensuring admissibility. This new metric makes another use of the drivable-area map, without requiring multiple annotations of trajectory futures. We couple this new metric with standard metrics from prior art, such as Average Displacement Error (ADE) and Final Displacement Error (FDE), to compare our model with existing baselines.

Additionally, we publish tools to replicate our data and results which we hope will advance the study of diverse trajectory forecasting. Our codes are available at: <https://github.com/kami93/CMU-DATF>.

## 2 Related Work

**Multimodal trajectory forecasting** requires a detailed understanding of the agent’s environment. Many works integrate information from multiple sensory cues, such as LiDAR point-cloud information to model the surrounding environment [17, 21, 22], high dimensional map data to model vehicle lane segmentation [3, 6, 33], and RGB image to capture the environmental context [17, 18, 24]. Other methods additionally fuse different combinations of interactions with the intention of jointly capturing all interactions between the agents [1, 11, 18, 24]. Without mechanisms to jointly and explicitly model such agent-to-scene and agent-to-agent relations, we hypothesize that these models are unable to capture complex interactions in the high-dimensional input space and propose methods to explicitly model these interactions via sophisticated attention mechanisms.

**Multi-agent modeling** aims to learn representations that summarize the behavior of one agent given its surrounding agents. These interactions are often modeled through either spatial-oriented methods or neural attention-based methods. Some of the spatial-oriented methods simply take into account agent-wise distances through a relative coordinate system [3, 13, 19, 22], while others utilize sophisticated pooling approaches across individual agent representations [8, 11, 17, 33]. On the other hand, the attention-based methods use the attention architecture to model multi-agent interaction in a variety of domains including pedestrians [9, 24, 28] and vehicles [18, 26]. In this paper, we employ the attention based cross-agent module to capture explicit agent-to-agent interactions. Even with the increasing number of agents around the ego-vehicle, our cross-agent module can successfully model the interactions between agents, as supported in one of our experiments.

**Diverse trajectory forecasting** involves stochastic modeling of future trajectories and sampling diverse predictions based on the model distribution. The Dynamic Bayesian network (DBN) is a common approach without deep generative models, utilized for modelling vehicle trajectories [10, 25, 30] and pedestrian actions [2, 15]. Although the DBN enables the models to consider physical process that generates agent trajectories, the performance is often limited for real traffic scenarios. Most state-of-the-art models utilize deep generative models such as GAN [11, 24, 33] and VAE [17] to encourage diverse predictions. However, these approaches mostly focus on generating multiple candidates while less focusing on analyzing the diversity across distributional modes. Recently, sophisticated sampling functions are proposed to tackle this issue, such as Diversity Sampling Function [32] and Latent Semantic Sampling [12]. Despite some promising empirical results, it remains difficult to evaluate both the diversity and admissibility of predictions. In this paper, we tackle the task of diverse trajectory forecasting with a special emphasis on admissibility in dynamic scenes and propose a new task metric that specifically assess models on the basis of these attributes.

## 3 Problem Formulation

We define the terminology that constitutes our problem. An *agent* is a dynamic on-road object that is represented as a sequence of 2D coordinates, i.e., a spatial position over time. We denote the position for agent  $a$  at time  $t$  as  $S_t^a \in \mathbb{R}^2$ ,

sequence of such positions from  $t_1$  to  $t_2$  as  $S_{t_1:t_2}^a$ , and the full sequence as (bold)  $\mathbf{S}^a$ . We set  $t = 0$  as *present*,  $t \leq 0$  as *past*, and  $t > 0$  as *prediction* or simply, *pred*. We often split the sequence into two parts, with respect to the *past* and *pred* subsequences: we denote these as  $\mathbf{S}_{\text{past}}^a$  and  $\mathbf{S}_{\text{pred}}^a$ . In order to clearly distinguish the predicted values from these variables, we use ‘hats’ such as  $\hat{S}_t^a$  and  $\hat{\mathbf{S}}_{\text{pred}}$ . A *scene* is a high-dimensional structured data that describes the present environmental context around the agent. For this, we utilize a bird’s eye view array, denoted as  $\Phi \in \mathbb{R}^{H \times W \times C}$ , where  $H$  and  $W$  are the sizes of field around the agent and  $C$  is the channel size of the scene, where each channel consists of distinct information such as the drivable area, position, and distance encodings.

Combining the *scene* and all *agent* trajectories yields an *episode*. In the combined setting, there are a variable number of agents which we denote using bold  $\mathbf{S} \equiv \{\mathbf{S}^1, \mathbf{S}^2, \dots, \mathbf{S}^A\}$  and as similarly to other variables, we may split it into two subsets,  $\mathbf{S}_{\text{past}}$  and  $\mathbf{S}_{\text{pred}}$  to represent the past and prediction segments. Since  $\mathbf{S}_{\text{past}}$  and  $\Phi$  serve as the observed information cue used for the prediction, they are often called *observation* simply being denoted as  $\mathcal{O} \equiv \{\mathbf{S}_{\text{past}}, \Phi\}$ .

We define *diversity* to be the level of coverage in a model’s predictions, across modes in a distribution representing all possible future trajectories. We denote the model distribution as  $q(\mathbf{S}_{\text{pred}}^a | \mathcal{O})$  and want the model to generate multiple samples interpreting each sample as an independent *hypothesis* that might have happened, given the same observation. We also acknowledge that encouraging a model’s predictions to be diverse, alone, is not sufficient for accurate and safe output; the model predictions should lie in the support of the true future trajectory distribution  $p(\mathbf{S}_{\text{pred}} | \mathcal{O})$ , i.e., predictions should be *admissible*. Given the observation  $\mathcal{O}$ , it is futile to predict samples around regions that are physically and statistically implausible to reach.

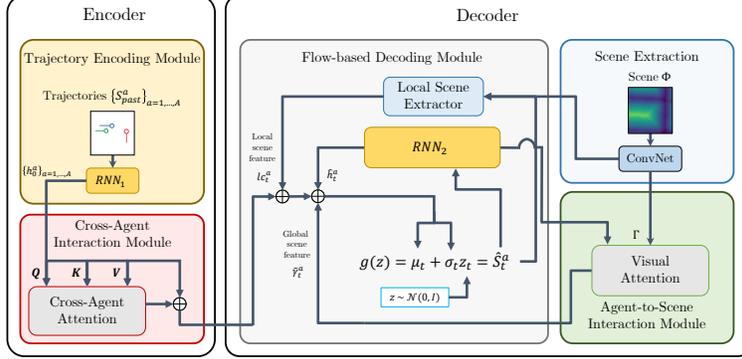
To summarize, this paper addresses the task of *diverse and admissible multi-agent trajectory forecasting*, by modeling multiple modes in the posterior distribution over the *prediction* trajectories, given the observation.

## 4 Proposed Approach

We hypothesize that future trajectories of human drivers should follow distributions of multiple modes, conditioned on the scene context and social behaviors of agents. Therefore, we design our model to explicitly capture both agent-to-scene interactions and agent-to-agent interactions with respect to each agent of interest. Through our objective function, we explicitly encourage the model to learn a distribution with multiple modes by taking into account past trajectories and attended scene context.

### 4.1 Model Architecture

As illustrated in Fig. 3, our model consists of an encoder-decoder architecture. The encoder includes the cross-agent interaction module. The decoder, on the other hand, comprises the agent-to-scene interaction module to capture the scene interactions. Please refer to Fig. 4 for a detailed illustration of our main proposed modules.



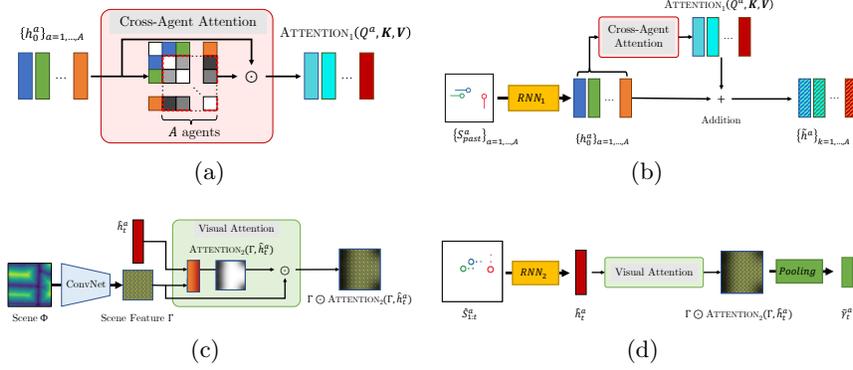
**Fig. 3.** Model Architecture. The model consists of an encoder-decoder architecture: the encoder takes as past agent trajectories and calculates cross-agent attention, and the flow-based decoder predicts future trajectories by attending scene contexts for each decoding step.

The *encoder* extracts past trajectory encoding for each agent, then calculates and fuses the interaction features among the agents. Given an observation, we encode each agent’s past trajectory  $S_{\text{past}}^a$  by feeding it to the trajectory encoding module. The module has the LSTM-based layer  $RNN_1$  to summarize the past trajectory. It iterates through the past trajectory with Eq. (1) and its final output  $h_0^a$  (at *present*  $t = 0$ ) is utilized as the agent embedding. The collection of the embeddings for all agents is then passed to the cross-agent interaction module, depicted in Fig. 4(a), which uses *self-attention* [27] to generate a cross-agent representation. We linearly transform each agent embedding to get a query-key-value triple,  $(Q^a, K^a, V^a)$ . Next, we calculate the interaction features through the self-attention layer  $\text{ATTENTION}_1$ . Finally, the fused agent encoding  $\tilde{h} \equiv \{\tilde{h}^1, \tilde{h}^2, \dots, \tilde{h}^A\}$  is calculated by adding each attended features to the corresponding embedding (see Eq. (2) and Fig. 4(b)).

$$h_t^a = \text{RNN}_1(S_{t-1}^a, h_{t-1}^a) \quad (1)$$

$$\tilde{h}^a = h_0^a + \text{ATTENTION}_1(Q^a, K, V) \quad (2)$$

The *decoder* takes the final encoding  $\tilde{h}^a$  and the scene context  $\Phi$  as inputs. We first extract the scene feature through convolutional neural networks i.e.,  $\Gamma = \text{CNN}(\Phi)$ . The decoder then autoregressively generates the future position  $\hat{S}_t^a$ , while referring to both local and global scene context from the agent-to-scene interaction module. The local scene feature is gathered using bilinear interpolation at the local region of  $\Gamma$ , corresponding to the physical position  $\hat{S}_{t-1}^a$ . Then, the feature is concatenated with  $\tilde{h}^a$  and processed thorough FC layers to form the local context  $lc_t^a$ . We call this part the local scene extractor. The global scene feature is calculated using *visual-attention* [31] to generate weighted scene features, as shown in Fig. 4(c). To calculate the attention, we first encode previous outputs  $\hat{S}_{1:t-1}^a$ , using a GRU-based  $RNN_2$  in Eq. (3), whose output is



**Fig. 4.** (a) Cross-agent attention. Interaction between each agent is modeled using attention, (b) Cross-agent interaction module. Agent trajectory encodings are corrected via cross-agent attention. (c) Visual attention. Agent-specific scene features are calculated using attention. (d) Agent-to-scene interaction module. Pooled vectors are retrieved from pooling layer after visual attention.

then used to calculate the pixel-wise attention  $\tilde{\gamma}_t^a$  at each decoding step, for each agent in Eq. (4).

$$\hat{h}_t^a = \text{RNN}_2(\hat{S}_{1:t-1}^a, \hat{h}_{t-1}^a) \quad (3)$$

$$\tilde{\gamma}_t^a = \text{POOL}(\mathbf{\Gamma} \odot \text{ATTENTION}_2(\hat{h}_t^a, \mathbf{\Gamma})) \quad (4)$$

The flow-based decoding module generates the future position  $\hat{S}_t^a$ . The module utilizes *Normalizing Flow* [20], a generative modeling method using bijective and differentiable functions; in particular, we choose an autoregressive design [14, 21, 22]. We concatenate  $\tilde{\gamma}_t^a$ ,  $\hat{h}_t^a$ , and  $lc_t^a$  then project them down to 6-dimensional vector using FC layers. We split this vector to  $\hat{\mu}_t \in \mathbb{R}^2$  and  $\hat{\sigma}_t \in \mathbb{R}^{2 \times 2}$ . Next, we transform a standard Gaussian sample  $z_t \sim \mathcal{N}(\mathbf{0}, I) \in \mathbb{R}^2$ , by the mapping  $g_\theta(z_t; \mu_t, \sigma_t) = \sigma_t \cdot z_t + \mu_t = \hat{S}_t^a$ , where  $\theta$  is the set of model parameters. ‘hats’ over  $\hat{\mu}_t$  and  $\hat{\sigma}_t$  are removed, in order to note that they went through the following details. To ensure the positive definiteness, we apply matrix exponential  $\sigma_t = \text{expm}(\hat{\sigma}_t)$  using the formula in [4]. Also, to improve the the physical *admissibility* of the prediction, we apply the constraint  $\mu_t = \hat{S}_{t-1}^a + \alpha(\hat{S}_{t-1}^a - \hat{S}_{t-2}^a) + \hat{\mu}_t$ , where  $\alpha$  is a model degradation coefficient. When  $\alpha = 1$ , the constraint is equivalent to *Verlet integration* [29], used in some previous works [21, 22], which gives the a perfect constant velocity (CV) prior to the model. However, we found empirically that, the model easily overfits to datasets when the the perfect CV prior is used, and perturbing the model with  $\alpha$  prevents overfitting; we choose  $\alpha = 0.5$ .

Iterating through time, we get the predictive trajectory  $\hat{S}_{\text{pred}}^a$  for each agent. By sampling multiple instances of  $z_{\text{pred}}$  and mapping them to trajectories, we get various hypotheses of future. Further details of the our network architecture and experiments for the degradation coefficient  $\alpha$  are given in the supplementary.

## 4.2 Learning

Our model learns to predict the distribution over the future trajectories of agents present in a given episode. In detail, we focus on predicting the conditional distribution  $p(\mathbf{S}_{\text{pred}}|\mathcal{O})$  where the future trajectory  $\mathbf{S}_{\text{pred}}$  depends on the observation. As described in the previous sections, our model incorporates a bijective and differentiable mapping between standard Gaussian prior  $q_0$  and the future trajectory distribution  $q_\theta$ . Such technique, commonly aliasd ‘normalizing flow’, enables our model not only to generate multiple candidate samples of future, but also to evaluate the ground-truth trajectory with respect to the predicted distribution  $q_\theta$  by using the inverse and the change-of-variable formula in Eq. (5).

$$q_\theta(S_{\text{pred}}^a|\mathcal{O}) = q_0(g^{-1}(S_{\text{pred}}^a)) \left| \det(\partial S_{\text{pred}}^a / \partial (g^{-1}(S_{\text{pred}}^a))) \right|^{-1} \quad (5)$$

As a result, our model can simply close the discrepancy between the predictive distribution  $q_\theta$  and the real world distribution  $p$  by optimizing our model parameter  $\theta$ . In particular, we choose to minimize the combination of forward and reverse cross-entropies, also known as ‘symmetric cross-entropy’, between the two distributions in Eq. (6). Minimizing the symmetric cross-entropy allows model to learn generating diverse and plausible trajectory, which is mainly used in [21].

$$\min_{\theta} H(p, q_\theta) + \beta H(q_\theta, p) \quad (6)$$

To realize this, we gather the ground-truth trajectories  $\mathbf{S}$  and scene context  $\Phi$  from the dataset that we assume to well reflect the real distribution  $p$ , then optimize the model parameter  $\theta$  such that 1) the density of the ground-truth future trajectories on top of the predicted distribution  $q_\theta$  is maximized and 2) the density of the predicted samples on top of the real distribution  $p$  is also maximized as described in Eq. (7). Since this objective is fully differentiable with respect to the model parameter  $\theta$ , we train our model using the stochastic gradient descent.

$$\min_{\theta} \mathbb{E}_{\mathbf{S}_{\text{pred}}, \mathcal{O} \sim p} \left[ \mathbb{E}_{S_t^a \in \mathbf{S}_{\text{pred}}} - \log q_\theta(S_t^a|\mathcal{O}) + \beta \mathbb{E}_{\hat{S}_t^a \sim q_\theta} - \log p(\hat{S}_t^a|\mathcal{O}) \right] \quad (7)$$

Such symmetric combination of the two cross-entropies guides our model to predict  $q_\theta$  that covers all plausible modes in the future trajectory while penalizing the bad samples that are less likely under the real distribution  $p$ . However, one major problem inherent in this setting is that we cannot actually evaluate  $p$  in practice. In this paper, we propose a new method which approximates  $p$  using the drivable-area map and we discuss it in the following subsection 4.3. Other optimization details are included in the supplementary.

## 4.3 The Drivable-area Map and Approximating $p$

We generate a binary mask that denote the drivable-area around the agents. We refer to this feature as the *drivable-area map* and utilize it for three different purposes: 1) deriving the approximated true trajectory distribution  $\tilde{p}$ , 2) calculating the diversity and admissibility metrics, and 3) building the scene input  $\Phi$ .

Particularly,  $\tilde{p}$  is a key component in our training objective, Eq. (6). Since the reverse cross-entropy penalizes the predicted trajectories with respect to the real distribution, the approximation should not underestimate some region of the real distribution, or diversity in the prediction could be discouraged. Previous works on  $\tilde{p}$  utilize the ground-truth (GT) trajectories to model it [21]. However, there is often only one GT annotation available thus deriving  $\tilde{p}$  based on the GT could assign awkwardly low density around certain region. To cope with such problem, our method assumes that every drivable locations are equally probable for future trajectories to appear in and that the non-drivable locations are increasingly less probable, proportional to the distance from the drivable-area. To derive it, we encode the distance on each non-drivable location using the distance transform on the drivable-area maps, then apply softmax over the transformed map to constitute it as a probability distribution. The visualizations of the  $\tilde{p}$  are available in Fig. 6. Further details on deriving  $\tilde{p}$  and the scene context input  $\Phi$ , as well as additional visualizations and qualitative results are given in the supplementary.

## 5 Experimental Setup

The primary goal in the following experiments is to evaluate our model, baselines, and ablations on the following criteria—(i) Leveraging mechanisms that explicitly model agent-to-agent and agent-to-scene interactions (experiments 1 and 2). (ii) Producing diverse trajectory predictions, while obeying admissibility constraints on the trajectory candidates, given different approximation methods for the true trajectory distribution  $p$  (experiment 3). (iii) Remaining robust to an increasing number of agents in the scene (agent complexity; experiment 4). (iv) Generalizing to other domains (experiment 5).

### 5.1 Dataset

We utilize two real world datasets to evaluate our model and the baselines: NUSCENES tracking [5] and ARGOVERSE motion forecasting [7]. NUSCENES contains 850 different real-world driving scenarios, where each spanning 20 seconds of frames and 3D box annotations for the surrounding objects. It also provides drivable-area maps. Based on this setting, we generate trajectories by associating the box annotations of the same agents. While NUSCENES provides trajectories for realistic autonomous driving scenarios, the number of episodes is limited around 25K. On the other hand, ARGOVERSE provides around 320K episodes from the real world driving along with the drivable-area maps. ARGOVERSE presents independent episodes spanning only 5s (2s for the past and 3s for the prediction), rather than providing long, continuing scenarios as in NUSCENES. However, this setting suffices to test our method and we evaluate baselines and our models on these two real-world datasets to provide complementary validations of each model’s diversity, admissibility, and generalizing performance across domains. Further details in the data processing are available in the supplementary.

### 5.2 Baseline Models

**Deterministic baselines:** We compare three deterministic models with our approach, to examine our model’s ability to capture agent-to-agent interaction:

*LSTM-based encoder-decoder* (LSTM), *LSTM with convolutional social pooling* [8] (CSP), and a deterministic version of *multi-agent tensor fusion* (MATF-D) [33]. For our deterministic model, we use an LSTM with our cross-agent attention module in the encoder, which we refer to as the *cross-agent attention model* (CAM). Because each model is predicated on an LSTM component, we set the capacity to be the same in all cases, to ensure fair comparison.

**Stochastic baselines:** We experiment three stochastic baselines. Our first stochastic baseline is a model based on a Variational Autoencoder structure, (DESIRE) [17], which utilizes scene contexts and an iterative refinement process. The second baseline model is a Generative Adversarial Network version of *multi-agent tensor fusion* (MATF-GAN) [33]. Our third baseline is the Reparameterized Pushforward Policy (R2P2-MA) [22] which is a modified version of R2P2 [21] for multi-agent prediction. To validate our model’s ability to extract scene information and generate diverse trajectories, multiple versions of our models are tested. While these models can be used as stand-alone models to predict diverse trajectories, comparison amongst these new models is equivalent to an ablation study of our final model. CAM-NF is a CAM model with a flow-based decoder. LOCAL-CAM-NF is CAM-NF with local scene features. GLOBAL-CAM-NF is LOCAL-CAM-NF with global scene features. Finally, ATTGLOBAL-CAM-NF is GLOBAL-CAM-NF with agent-to-scene attention, which is our main proposed model.

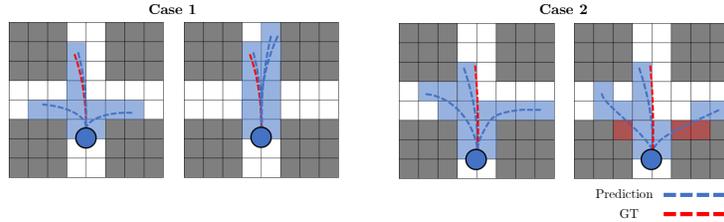
### 5.3 Metrics

We define multiple metrics that provide a thorough interpretation about the behavior of each model in terms of precision, diversity, and admissibility. To evaluate precision, we calculate Euclidean errors: ADE (*average displacement error*) and FDE (*final displacement error*), or ERROR to denote both. To evaluate multiple hypotheses, we use the average and the minimum ERROR among  $K$  hypotheses:  $\text{AVGERROR} = \frac{1}{k} \sum_{i=1}^k \text{ERROR}^{(i)}$  and  $\text{MINERROR} = \min\{\text{ERROR}^{(1)}, \dots, \text{ERROR}^{(k)}\}$ . A large AVGERROR implies that predictions are spread out, and a small MINERROR implies at least one of predictions has high precision. From this observation, we define a new evaluation metric that capture diversity in predictions: the ratio of AVGFDE to MINFDE, namely RF. RF is robust to the variability of magnitude in velocity in predictions hence provides a handy tool that can distinguish between predictions with multiple modes (diversity) and predictions with a single mode (perturbation).

$$\text{Ratio of AVGFDE to MINFDE (RF)} = \frac{\text{AVGFDE}}{\text{MINFDE}} \quad (8)$$

$$\text{Drivable Area Occupancy (DAO)} = \frac{\text{count}(\text{traj}_{\text{pix}})}{\text{count}(\text{driv}_{\text{pix}})} \quad (9)$$

We also report performance on additional metrics that are designed to capture diversity and admissibility in predictions. While RF measures the spread of predictions in Euclidean distance, DAO measures diversity in predictions that are only admissible, and DAC measures extreme off-road predictions that defy admissibility. We follow [7] in the use of *Drivable Area Count* (DAC),  $\text{DAC} =$



**Fig. 5.** We motivate the need for multiple metrics, to assess diversity and admissibility. Case 1: DAO measures are equal, even though predictions have differing regard for the modes in the posterior distribution. Case 2: rF measures are equal, despite differing regard for the cost of leaving the drivable area. In both cases, it is important to distinguish between conditions—we do this by using DAO, rF, and DAC together.

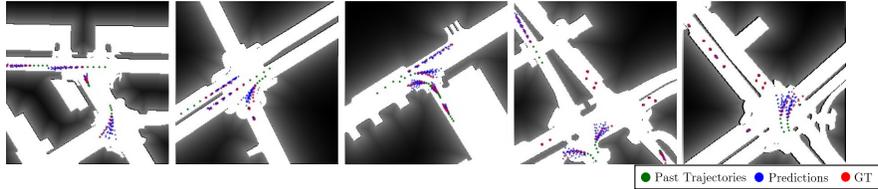
$\frac{k-m}{k}$ , where  $m$  is the number of predictions that go out of the drivable area and  $k$  is the number of hypotheses per agent. Next, we propose a new metric, *Drivable Area Occupancy* (DAO), which measures the proportion of pixels that predicted trajectories occupy in the drivable-area. Shown in Eq. (9),  $\text{count}(\text{traj}_{\text{pix}})$  is the number of pixels occupied by predictions and  $\text{count}(\text{driv}_{\text{pix}})$  is the total number of pixels of the drivable area, both within a pre-defined grid around the ego-vehicle.

DAO may seem a reasonable standalone measure of capturing both diversity and admissibility, as it considers diversity in a reasonable region of interest. However, DAO itself cannot distinguish between *diversity* (Section 3) and arbitrary stochasticity in predictions, as illustrated by Case 1 in Fig. 5: although DAO measures of both predictions are equal, the causality behind each prediction is different and we must distinguish the two. rF and DAO work in a complementary way and we, therefore, use both for measuring diversity. To assure the *admissibility* of predictions, we use DAC which explicitly counts off-road predictions, as shown by Case 2 in Fig. 5. As a result, assessing predictions using DAO along with rF and DAC provides a holistic view of the quantity and the quality of diversity in predictions.

For our experiments, we use MINADE and MINFDE to measure *precision*, and use rF, DAC, and DAO to measure both *diversity* and *admissibility*. Due to the nature of DAO, where the denominator in our case is the number of overlapping pixels in a  $224 \times 224$  grid, we normalize it by multiplying by 10,000 when reporting results. For the multi-agent experiment (shown in Table 4), relative improvement (RI) is calculated as we are interested in the relative improvement as the number of agents increases. For all metrics, the number of trajectory hypotheses should be set equally for fair comparison of models. If not specified, the number of hypotheses  $k$  is set to 12 when reporting the performance metrics.

## 6 Results and Discussion

In this section, we show experimental results on numerous settings including the comparison with the baseline, and ablation studies of our model. We first show the effect of our cross-agent interaction module and agent-to-scene interaction



**Fig. 6.** Our map loss and corresponding model predictions. Each pixel on our map loss denotes probability of future trajectories; higher probability values are represented by brighter pixels. Our approach generates diverse and admissible future trajectories. More visualizations of qualitative results are provided in the supplementary.

module on the model performance, then we analyze the performance with respect to different numbers of agents, and other datasets. All experiments are measured with MINADE, MINFDE, RF, DAC, and DAO for holistic interpretation.

**Table 1.** Deterministic models on NUSCENES. Our proposed model outperforms the existing baselines.

Model	MINADE ( $\downarrow$ )	MINFDE ( $\downarrow$ )
LSTM	1.186	2.408
CSP	1.390	2.676
MATF-D [33]	1.261	2.538
CAM (OURS)	<b>1.124</b>	<b>2.318</b>

**Table 2.** Stochastic models on NUSCENES. Our models outperform the existing baselines, achieving the best precisions, diversity, and admissibility. Improvements indicated by arrows.

Model	MINADE ( $\downarrow$ )	MINFDE ( $\downarrow$ )	RF ( $\uparrow$ )	DAO ( $\uparrow$ )	DAC ( $\uparrow$ )
DESIRE [17]	1.079	1.844	1.717	16.29	0.776
MATF-GAN [33]	1.053	2.126	1.194	11.64	0.910
R2P2-MA [21]	1.179	2.194	1.636	<b>25.65</b>	0.893
CAM-NF (OURS)	0.756	1.381	2.123	23.15	0.914
LOCAL-CAM-NF (OURS)	0.774	1.408	2.063	22.58	<b>0.921</b>
GLOBAL-CAM-NF (OURS)	0.743	1.357	2.106	22.65	<b>0.921</b>
ATTGLOBAL-CAM-NF (OURS)	<b>0.639</b>	<b>1.171</b>	<b>2.558</b>	22.62	0.918

**Effectiveness of cross-agent interaction module:** We show the performance of one of our proposed models CAM, which utilizes our cross-agent attention module, along with three deterministic baselines as shown in Table 1. CSP models the interaction through layers of convolutional networks, and the interaction is implicitly calculated within the receptive field of convolutional layers. MATF-D is an extension of convolutional social pooling with scene information. CAM explicitly defines the interaction between each agent by using attention. The

**Table 3.** Training ATTGLOBAL-CAM-NF using the proposed annotation-free  $\tilde{p}$  outperforms the annotation-dependent counterpart (MSE) in NUSCENES.

Model	MINADE ( $\downarrow$ )	MINFDE ( $\downarrow$ )	rF ( $\uparrow$ )	DAO ( $\uparrow$ )	DAC ( $\uparrow$ )
ATTGLOBAL-CAM-NF (MSE)	0.735	1.379	1.918	21.48	<b>0.924</b>
ATTGLOBAL-CAM-NF	<b>0.638</b>	<b>1.171</b>	<b>2.558</b>	<b>22.62</b>	0.918

**Table 4.** Multi-agent experiments on NUSCENES (MINFDE). RI denotes relative improvements of MINFDE between 10 and 1 agent. Our approach best models multi-agent interactions.

Model	1 agent	3 agents	5 agents	10 agents	RI(1-10)
LSTM	2.736	2.477	2.442	2.268	17.1%
CSP	2.871	2.679	2.671	2.569	10.5%
DESIRE [17]	2.150	1.846	1.878	1.784	17.0%
MATF GAN [33]	2.377	2.168	2.150	2.011	15.4%
R2P2-MA [21]	2.227	2.135	2.142	2.048	8.0%
ATTGLOBAL-CAM-NF (OURS)	<b>1.278</b>	<b>1.158</b>	<b>1.100</b>	<b>0.964</b>	<b>24.6%</b>

result shows that CAM outperforms other baselines in both MINADE and MINFDE, indicating that the explicit way of modeling agent-to-agent interaction performs better in terms of precision than an implicit way of modeling interaction using convolutional networks used in CSP and MATF-D. Interestingly, CAM outperforms MATF-D that utilizes scene information. This infers that the cross-agent interaction module has ability to learn the structure of roads and permissible region given the trajectories of surrounding agents.

**Effectiveness of agent-to-scene interaction module:** The performance of stochastic models is compared in Table 2. We experiment with removing scene processing operations in the decoder to validate the importance of our proposed agent-to-scene interaction module. As mentioned previously, generating multiple modes of sample requires a strong scene processing module and a diversity-oriented decoder. Our models achieves the best precision. MATF-GAN has a small rF showing that predictions are unimodal, while other models such as VAE-based model (DESIRE) and flow-based models (R2P2-MA and OURS) show spread in their predictions. R2P2-MA shows the highest DAO. We note that OURS has a comparable DAO while keeping the highest rF and DAC, indicating that our models exhibit diverse and admissible predictions by accurately utilizing scene context.

**Effectiveness of new  $\tilde{p}$ :** We compare two different versions of our ATTGLOBAL-CAM-NF. One is trained using mean squared error (MSE) between  $\mathbf{S}_{pred}$  and  $\hat{\mathbf{S}}_{pred}$  as an example of annotation-based approximation for  $p$ , while the other is trained with our drivable area-based (annotation-free) approximation of  $p$  in Table 3. Using our new approximation in training shows superior results in most of the reported metrics. In particular, the precision and the diversity of predictions increases drastically as reflected in MINERROR, DAO, and rF while DAC remains comparable. Thus our  $\tilde{p}$  considers admissibility while improving precision and diversity i.e drivable-area related approximate enhances the estimate on additional trajectories over the most probable one.

**Complexity from number of agents:** We experiment with varying number of surrounding agents as shown in Table 4. Throughout all models, the perfor-

**Table 5.** Results of baseline models (upper partition) and our proposed models (lower partition). ATTGLOBAL-CAM-NF is our full proposed model and others in the lower partition are the ablations. The metrics are abbreviated as follows: MINADE(**A**), MINFDE(**B**), RF(**C**), DAO(**D**), DAC(**E**). Improvements indicated by arrows.

Model	ARGOVERSE					NUSCENES				
	A (↓)	B (↓)	C (↑)*	D (↑)*	E (↑)*	A (↓)	B (↓)	C (↑)*	D (↑)*	E (↑)*
LSTM	1.441	2.780	1.000	3.435	0.959	1.186	2.408	1.000	3.187	0.912
CSP	1.385	2.567	1.000	3.453	0.963	1.390	2.676	1.000	3.228	0.900
MATF-D [33]	1.344	2.484	1.000	3.372	0.965	1.261	2.538	1.000	3.191	0.906
DESIRE [17]	0.896	1.453	3.188	15.17	0.457	1.079	1.844	1.717	16.29	0.776
MATF-GAN [33]	1.261	2.313	1.175	11.47	0.960	1.053	2.126	1.194	11.64	0.910
R2P2-MA [22]	1.108	1.771	3.054	<b>37.18</b>	0.955	1.179	2.194	1.636	<b>25.65</b>	0.893
CAM	1.131	2.504	1.000	3.244	<b>0.973</b>	1.124	2.318	1.000	3.121	<b>0.924</b>
CAM-NF	0.851	1.349	2.915	32.89	0.951	0.756	1.381	2.123	23.15	0.914
LOCAL-CAM-NF	0.808	1.253	3.025	31.80	0.965	0.774	1.408	2.063	22.58	0.921
GLOBAL-CAM-NF	0.806	1.252	3.040	31.59	0.965	0.743	1.357	2.106	22.65	0.921
ATTGLOBAL-CAM-NF	<b>0.730</b>	<b>1.124</b>	<b>3.282</b>	28.64	0.968	<b>0.639</b>	<b>1.171</b>	<b>2.558</b>	22.62	0.918

mance increases as the number of agents increases even though we observe that many agents in the surrounding do not move significantly. In terms of relative improvement RI, as calculated between 1 agent and 10 agents, our model has the most improvement, indicating that our model makes the most use of the fine-grained trajectories of surrounding agents to generate future trajectories.

**Generalizability across datasets:** We compare our model with baselines extensively across another real world dataset ARGOVERSE to test generalization to different environments. We show results in Table 5 where we outperform or achieve comparable results as compared to the baselines. For ARGOVERSE, we additionally outperform MFP3 [26] in MINFDE with 6 hypotheses: our full model shows a MINFDE of 0.915, while MFP3 achieves 1.399.

## 7 Conclusion

In this paper, we tackled the problem of generating diverse and admissible trajectory predictions by understanding each agent’s multimodal context. We proposed a model that learns agent-to-agent interactions and agent-to-scene interactions using attention mechanisms, resulting in better prediction in terms of precision, diversity, and admissibility. We also developed a new approximation method that provides richer information about the true trajectory distribution and allows more accurate training of flow-based generative models. Finally, we present new metrics that provide a holistic view of the quantity and the quality of diversity in prediction, and a NUSCENES trajectory extraction code to support future research in diverse and admissible trajectory forecasting.

**Acknowledgements.** This work was supported in part by the Technology Innovation Program under Grant 10083646 (Development of Deep Learning-Based Future Prediction and Risk Assessment Technology considering Inter-vehicular Interaction in Cut-in Scenario), funded by the Ministry of Trade, Industry, and Energy, South Korea. We also acknowledge the anonymous reviewers for their constructive 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: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 961–971 (2016)
2. Ballan, L., Castaldo, F., Alahi, A., Palmieri, F., Savarese, S.: Knowledge transfer for scene-specific motion prediction. In: European Conference on Computer Vision. pp. 697–713. Springer (2016)
3. Bansal, M., Krizhevsky, A., Ogale, A.: Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. arXiv preprint arXiv:1812.03079 (2018)
4. Bernstein, D.S., So, W.: Some explicit formulas for the matrix exponential. *IEEE Transactions on Automatic Control* **38**(8), 1228–1232 (1993)
5. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11621–11631 (2020)
6. Casas, S., Luo, W., Urtasun, R.: Intentnet: Learning to predict intention from raw sensor data. In: Conference on Robot Learning. pp. 947–956 (2018)
7. Chang, M.F., Lambert, J., Sangkloy, P., Singh, J., Bak, S., Hartnett, A., Wang, D., Carr, P., Lucey, S., Ramanan, D., et al.: Argoverse: 3d tracking and forecasting with rich maps. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8748–8757 (2019)
8. Deo, N., Trivedi, M.M.: Convolutional social pooling for vehicle trajectory prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 1468–1476 (2018)
9. Fernando, T., Denman, S., Sridharan, S., Fookes, C.: Soft+ hardwired attention: An lstm framework for human trajectory prediction and abnormal event detection. *Neural networks* **108**, 466–478 (2018)
10. Gindele, T., Brechtel, S., Dillmann, R.: Learning driver behavior models from traffic observations for decision making and planning. *IEEE Intelligent Transportation Systems Magazine* **7**(1), 69–79 (2015)
11. 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)
12. Huang, X., McGill, S., DeCastro, J., Fletcher, L., Leonard, J., Williams, B., Rosman, G.: Diversitygan: Diversity-aware vehicle motion prediction via latent semantic sampling. *IEEE Robotics and Automation Letters* (2020)
13. Kim, B., Kang, C.M., Kim, J., Lee, S.H., Chung, C.C., Choi, J.W.: Probabilistic vehicle trajectory prediction over occupancy grid map via recurrent neural network. In: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC). pp. 399–404. IEEE (2017)
14. Kingma, D.P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., Welling, M.: Improved variational inference with inverse autoregressive flow. In: Advances in neural information processing systems. pp. 4743–4751 (2016)
15. Kooij, J.F.P., Schneider, N., Flohr, F., Gavrilu, D.M.: Context-based pedestrian path prediction. In: European Conference on Computer Vision. pp. 618–633. Springer (2014)
16. Krajewski, R., Bock, J., Kloeker, L., Eckstein, L.: The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC). pp. 2118–2125. IEEE (2018)

17. 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)
18. Ma, Y., Zhu, X., Zhang, S., Yang, R., Wang, W., Manocha, D.: Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 6120–6127 (2019)
19. Park, S.H., Kim, B., Kang, C.M., Chung, C.C., Choi, J.W.: Sequence-to-sequence prediction of vehicle trajectory via lstm encoder-decoder architecture. In: 2018 IEEE Intelligent Vehicles Symposium (IV). pp. 1672–1678. IEEE (2018)
20. Rezende, D.J., Mohamed, S.: Variational inference with normalizing flows. arXiv preprint arXiv:1505.05770 (2015)
21. Rhinehart, N., Kitani, K.M., Vernaza, P.: R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 772–788 (2018)
22. Rhinehart, N., McAllister, R., Kitani, K., Levine, S.: Precog: Prediction conditioned on goals in visual multi-agent settings. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2821–2830 (2019)
23. Rudenko, A., Palmieri, L., Herman, M., Kitani, K.M., Gavrilu, D.M., Arras, K.O.: Human motion trajectory prediction: A survey. *The International Journal of Robotics Research* p. 0278364920917446 (2019)
24. Sadeghian, A., Kosaraju, V., Sadeghian, A., Hirose, N., Rezafofighi, H., Savarese, S.: Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1349–1358 (2019)
25. Schulz, J., Hubmann, C., Löchner, J., Burschka, D.: Interaction-aware probabilistic behavior prediction in urban environments. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3999–4006. IEEE (2018)
26. Tang, C., Salakhutdinov, R.R.: Multiple futures prediction. In: Advances in Neural Information Processing Systems. pp. 15398–15408 (2019)
27. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
28. 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)
29. Verlet, L.: Computer” experiments” on classical fluids. i. thermodynamical properties of lennard-jones molecules. *Physical review* **159**(1), 98 (1967)
30. Xie, G., Gao, H., Qian, L., Huang, B., Li, K., Wang, J.: Vehicle trajectory prediction by integrating physics-and maneuver-based approaches using interactive multiple models. *IEEE Transactions on Industrial Electronics* **65**(7), 5999–6008 (2017)
31. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: International conference on machine learning. pp. 2048–2057 (2015)
32. Yuan, Y., Kitani, K.M.: Diverse trajectory forecasting with determinantal point processes. In: International Conference on Learning Representations (2020)
33. Zhao, T., Xu, Y., Monfort, M., Choi, W., Baker, C., Zhao, Y., Wang, Y., Wu, Y.N.: Multi-agent tensor fusion for contextual trajectory prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 12126–12134 (2019)