

# RAIN: Reinforced Hybrid Attention Inference Network for Motion Forecasting

Jiachen Li<sup>1,2\*</sup> Fan Yang<sup>3</sup> Hengbo Ma<sup>2</sup> Srikanth Malla<sup>1</sup> Masayoshi Tomizuka<sup>2</sup> Chiho Choi<sup>1</sup>

<sup>1</sup>Honda Research Institute USA

<sup>2</sup>University of California, Berkeley

<sup>3</sup>Carnegie Mellon University

{jiachen.li, hengbo\_ma, tomizuka}@berkeley.edu fanyang3@andrew.cmu.edu

{smalla, cchoi}@honda-ri.com

## Abstract

Motion forecasting plays a significant role in various domains (e.g., autonomous driving, human-robot interaction), which aims to predict future motion sequences given a set of historical observations. However, the observed elements may be of different levels of importance. Some information may be irrelevant or even distracting to the forecasting in certain situations. To address this issue, we propose a generic motion forecasting framework (named RAIN) with dynamic key information selection and ranking based on a hybrid attention mechanism. The general framework is instantiated to handle multi-agent trajectory prediction and human motion forecasting tasks, respectively. In the former task, the model learns to recognize the relations between agents with a graph representation and to determine their relative significance. In the latter task, the model learns to capture the temporal proximity and dependency in long-term motions. We also propose an effective double-stage training pipeline with an alternating training strategy to optimize the parameters in different modules of the framework. We validate the framework on both synthetic simulations and motion forecasting benchmarks in different domains, demonstrating that our method not only achieves state-of-the-art forecasting performance, but also provides interpretable and reasonable hybrid attention weights.

## 1. Introduction

Motion forecasting has been widely studied in various domains, such as physical systems, human skeletons, and multi-agent interacting systems (e.g., traffic participants, sports players, etc). The problem is formulated as to predict future states or trajectories based on historical spatio-temporal observations. However, the observed information may be of different levels of significance and in some situations not all the information is relevant for the forecasting.

\*Work done during Jiachen’s internship at Honda Research Institute.

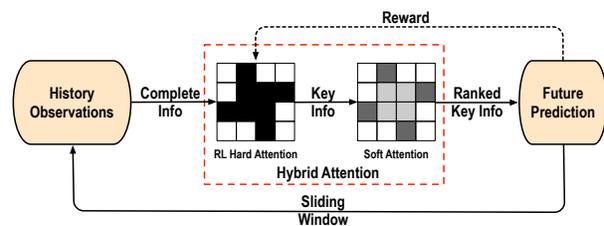


Figure 1. A high-level diagram of the proposed general prediction framework with key information/element selection and ranking, which consists of two major parts: an RL-based hard attention mechanism to discriminate key information from complete observations and a soft attention mechanism to further figure out relative significance of the key information. The whole procedure can be iteratively applied over time with a sliding window to enable dynamic selection of key information to adapt to evolving situations.

Moreover, the key information may be varying as the situation evolves, which motivates the forecasting approach to dynamically adjust its attention to different subsets of observations. Here we provide two illustrative real-world examples where key information is naturally selected based on either spatial relations or temporal dependencies. An on-road vehicle usually only needs to pay attention to the traffic participants that are interacting or having a conflict with itself, so only a subset of observations are indeed relevant when predicting its future behavior. For human motion forecasting, it is observed that humans tend to repeat their motions, which motivates dynamic attention to different segments of previous motions given the current observation.

Attention mechanisms have been widely adopted to learn the relative importance of elements. There are two major types of attention mechanisms in literature: soft attention and hard attention [63]. The soft attention is usually performed by applying a score function to input features followed by a softmax function to obtain the attention weights in the range of  $[0, 1]$ . These operations are fully differentiable which can be trained by back-propagation with typical gradient-based optimizers. However, the softmax func-

tion tends to assign non-zero attention weights to irrelevant or unimportant elements, which dilutes the attention given to the truly significant information [47, 54]. In contrast, the hard attention mechanism can force the model to only pay attention to the relevant information while discard the others entirely to reduce information redundancy. The hard attention weights can be only binary: 0 (discarded) or 1 (retained). However, the hard attention is not differentiable due to the  $\text{argmax}$  operation, which needs to be optimized by reinforcement learning (RL) algorithms (e.g., deep Q-learning [41], policy gradient [60]). Recently, some alternatives to the traditional hard attention have been proposed based on the approximation of evidence lower bound (ELBO) [38], which can be trained end-to-end.

The general idea of selecting the most important information/elements with hard attention has been applied to several different domains, such as computer vision and natural language processing. Wu et al [62] introduced a key frame selection framework based on multi-agent reinforcement learning for video based human activity recognition. Wang et al [55] presented a framework for informative view selection from multiple indoor cameras to recognize human actions. Gao et al [14] proposed a hard and channel-wise attention network for graph representation learning. Shen et al [47] illustrated a reinforced self-attention network to figure out sparse dependencies between tokens in a sentence.

However, the efficacy of hard attention in motion forecasting tasks, to the best of our knowledge, still remains largely unexplored so far. Besides, many existing works including the aforementioned ones, pre-define a fixed number of elements to pay attention to, which may be unsatisfactory in the scenarios where the amount of key information/elements is varying. For example, the motion of a certain entity in an interacting system may be affected by a varying number of entities at different time, thus a fixed number of selected elements may be redundant or insufficient in different situations. To address this issue, we propose a reinforcement learning based hard attention mechanism for motion forecasting, which does not enforce any constraints on the amount of key elements. It is even possible in some situations that no element or all the elements are selected based on their significance. In the multi-agent setting with a graph representation, learning hard attention can also be treated and interpreted as graph structure/topology learning [12, 32]. Since the selected key information may be still at different levels of importance, we propose to employ soft attention as a ranking mechanism to further discriminate relative importance.

To the best of our knowledge, we are the first to propose a hybrid attention based framework for motion forecasting, which is illustrated in Figure 1. The main contributions of this paper are summarized as:

- We propose a general motion forecasting framework

(named RAIN) with dynamic key information/element selection and ranking via a hybrid attention mechanism.

- We propose an effective double-stage training pipeline with an alternating training strategy to improve different modules in the framework alternatively.

- We instantiate the general framework and propose a novel graph-based model for multi-agent trajectory forecasting. We also demonstrate the general idea on human skeleton motion forecasting, where a state-of-the-art model [59] is employed as a part of our framework. We validate the proposed framework on both domains and our method achieves the state-of-the-art performance consistently.

## 2. Related Work

### 2.1. Trajectory/Motion Forecasting

Many research efforts have been devoted to motion forecasting in various domains. Here we particularly provide a brief review of literature on physical systems, highly interactive traffic scenarios and skeleton based human motions, which are closely related to this work.

- **Physical systems:** Learning the dynamics of physical systems involving multiple interacting elements have been studied in some recent works, either from simulated trajectories [3, 19, 23, 30] or from generated videos [25, 52, 58], where graph neural network is used to model interactions. Some of them assume a known graph topology based on prior knowledge, while the others infer the underlying structure explicitly or implicitly. Our approach falls into the second category, where we propose a novel fashion of relational inference based on reinforcement learning.

- **Traffic scenarios:** A number of works attempt to forecast future trajectories or behaviors of heterogeneous traffic participants (e.g., pedestrians, vehicles, cyclists) [5, 8, 13, 20, 24, 28, 35–37, 44, 61]. In order to model the relations or interactions between different entities, many information aggregation techniques have been developed, including but not limited to, feature concatenation [28], attention mechanisms [20, 24, 29, 54], social pooling [1, 9], message passing across graphs [7, 20, 30], etc. Moreover, recent works are also putting more emphasis on environmental modeling to leverage physical and semantic constraints such as road layout and traffic rules in driving scenarios [15, 44, 50].

- **Human motions:** Recurrent neural networks are widely utilized in human motion forecasting [11, 16, 22]. Recent works also proposed to adopt feed-forward methods for effective encoding of long-term motion history [27, 39]. Besides, attention mechanisms were also employed due to its flexibility and efficacy. In [51], Tang et al adopted a framework attention mechanism to summarize the pose history. In [59], Mao et al proposed a motion attention based method, which achieved the previous state-of-the-art performance. In this paper, we apply our hard attention module on top of

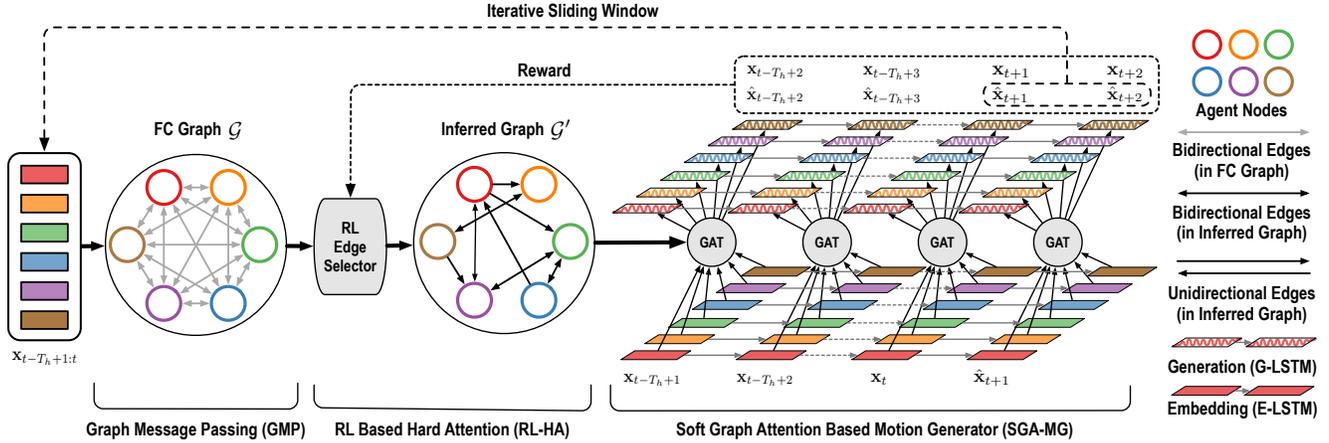


Figure 2. The diagram of the proposed hybrid attention based prediction framework for multi-agent interacting systems, which consists of three major components: a graph message passing module (GMP), a RL based hard attention module (RL-HA) and a soft graph attention based motion generator (SGA-MG). GMP aims to pass the information across the fully connected graph  $\mathcal{G}$ . RL-HA aims to figure out truly relevant information (edges) and output an inferred relation graph  $\mathcal{G}'$ , while SGA-MG aims to further rank the significance of the key information with soft attention and generate future motions with recurrent neural networks. The prediction process can be either one-shot (with static relation graph  $\mathcal{G}'$ ) or progressive with a sliding window (with dynamic relation graph  $\mathcal{G}'$ ). Best viewed in color.

their soft attention model [59] to illustrate the efficacy of our proposed hybrid attention based framework.

## 2.2. Attention Mechanisms

Attention mechanisms have been widely used in deep neural networks due to their efficacy and efficiency for sequence modeling and information fusion. Existing works have demonstrated the superiority of attention in a broad range of domains, such as natural language processing [43, 47, 57], image captioning [2, 56, 63], saliency detection [33, 64], time-series modeling [10, 45], human activity recognition [6, 48, 49]. Most of these papers either solely apply soft attention to obtain contextual embeddings or solely apply hard attention to select important elements. The only exception is [47] in which a hybrid attention was presented to figure out sparse dependencies between sequential tokens in a sentence. However, their method can only be applied to sequence modeling, which is hard to generalize to the motion forecasting task of multi-agent interacting systems. For the motion forecasting task, some existing works only adopt soft attention to weight the complete observed information [20, 24, 28, 29, 54, 59], while our method can discriminate the most relevant elements via the hybrid attention.

## 3. Problem Overview

Since the proposed approach can be widely applied to various domains, we introduce the problem formulation in a general way. Define a multivariate dynamic system

$$\mathbf{X}_{t+1:t+T_f} = f(\mathbf{X}_{t-T_h+1:t}, \mathbf{C}), \quad (1)$$

where  $\mathbf{X}_t = \{\mathbf{x}_t^i, i = 1, \dots, N\}$  denotes the system state at time  $t$  and  $\mathbf{C} = \{\mathbf{c}^i, i = 1, \dots, N\}$  denotes optional con-

text information or external factors.  $N$  is the total number of variables which have a specific meaning in different domains. The goal of this work is to approximate the conditional distribution  $p(\mathbf{X}_{t+1:t+T_f} | \mathbf{X}_{t-T_h+1:t}, \mathbf{C})$ , where  $T_h$  and  $T_f$  denote the history and prediction horizon.

For a multi-agent interacting system, the variables refer to the involved homogeneous/heterogeneous interacting entities, where the state may include position, velocity, etc. For a multivariate time series such as human motions, the variables refer to a set of human skeletons, where the state may include joint coordinates or relative angles.

In this paper, we instantiate and apply the proposed general prediction framework illustrated in Figure 1 to both multi-agent systems with spatially interacting agents (e.g., physical systems, traffic participants), and multivariate time series with temporal dependency (e.g., skeleton-based human motions). For long-term prediction of multi-agent interacting systems, we propose a novel, complete model architecture based on graph representation in Section 4. For multivariate time series (i.e., human motions), we present an effective way to build our proposed hard attention module on top of a state-of-the-art soft attention based model for human motion prediction [59] in Section 5.

## 4. RAIN for Multi-Agent Interacting System

### 4.1. Framework Overview

An illustrative diagram is shown in Figure 2 to introduce the pipeline and three major components: graph message passing module (GMP), RL based hard attention module (RL-HA) and soft graph attention based motion generator (SGA-MG), which cooperate closely to improve the final

prediction performance. More specifically, for the prediction of a certain target entity, GMP collects information from other entities across graph  $\mathcal{G}$ . RL-HA discriminates the key relevant elements from the complete observations and provides SGA-MG with an inferred relation graph  $\mathcal{G}'$  with only selected edges, which is a natural generalization of the traditional hard attention to graph representation. SGA-MG uses soft attention weights to rank the relative importance of key information and generates future trajectories. The prediction together with the ground truth provides rewards to RL-HA during the training phase to guide the improvement of the RL edge selector. GMP is pre-trained to collect contextual information across the whole graph. SGA-MG is pre-trained with a fully connected topology in order to improve training efficiency and stability as well as to enable informative initial reward.

## 4.2. Graph Message Passing

It is natural to represent a multi-agent system with  $N$  entities as a fully connected (FC) graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V} = \{\mathbf{v}_i, i = 1, \dots, N\}$  and  $\mathcal{E} = \{\mathbf{e}_{ij}, i, j = 1, \dots, N\}$ .  $\mathbf{v}_i$  denotes node  $i$ 's attribute and  $\mathbf{e}_{ij}$  denotes the edge attribute from sender node  $j$  to receiver node  $i$ . The node attribute consists of a self-attribute to store the individual information, a social-attribute to store other entities' information, and a context-attribute to include the agent's context information. More formally, we have  $\mathbf{v}_i^{\text{self}} = f_s^m(\mathbf{x}_{t-T_h+1:t}^i)$ ,  $\mathbf{v}_i^{\text{neighbor}} = f_n^m(\mathbf{x}_{t-T_h+1:t}^i)$ ,  $\mathbf{v}_i^{\text{context}} = f_c(\mathbf{c}^i)$ , where  $m \in \{1, \dots, M\}$ ,  $M$  is number of agent types.  $f_s^m$  and  $f_n^m$  are state embedding functions, and  $f_c$  is context embedding function. Different state embedding functions corresponding to certain agent types are applied to heterogeneous agents.

Since the relationship between a pair of agents is not only determined by their own behaviors but also affected by other agents in the scene, it is not sufficient to use only a pair of self node attributes to determine the existence of an edge. Therefore, we apply a round of message passing to collect the contextual information across the FC graph  $\mathcal{G}$  by

$$\alpha_{ij} = \frac{\exp\left(\text{MLP}\left(\left[\mathbf{v}_i^{\text{self}} \parallel \mathbf{v}_j^{\text{neighbor}}\right]\right)\right)}{\sum_{k \in \mathcal{N}_i} \exp\left(\text{MLP}\left(\left[\mathbf{v}_i^{\text{self}} \parallel \mathbf{v}_k^{\text{neighbor}}\right]\right)\right)}, \quad (2)$$

$$\mathbf{v}_i^{\text{social}} = f_v\left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbf{v}_j^{\text{neighbor}}\right), \quad (3)$$

where  $f_v$  is the social attribute update function and  $\parallel$  denotes the concatenation operation.  $\mathcal{N}_i$  denotes the set of one-hop neighbors of node  $i$ . MLP refers to multi-layer perceptron. The complete node attribute is

$$\mathbf{v}_i = f_{\text{enc}}\left(\left[\mathbf{v}_i^{\text{self}}, \mathbf{v}_i^{\text{social}}, \mathbf{v}_i^{\text{context}}\right]\right). \quad (4)$$

## 4.3. Hard Attention: Key Information Selection

As shown in Figure 1, the RL-based hard attention module serves as a key information selector, which takes the

complete history observations as input and discriminates the truly relevant information while totally discard the rest. More specifically, in the context of multi-agent system, the RL-HA module is expected to figure out truly influencing factors when predicting the motions of a certain agent. In other words, the goal is to assert the existence of each edge in the FC graph based on the observations so that the redundant information is discarded in the prediction.

The selection of key edges naturally fits into a reinforcement learning framework. The definition of observations, actions and reward functions are elaborated in the following, while the training procedures are left to Section 4.6.

**Observations:** The observation  $O$  of RL-agent at RL-step  $\eta$  ( $\leq T_{\text{RL}}$ ) includes a pair of node attributes  $\mathbf{v}_i$  and  $\mathbf{v}_j$  as well as the current edge selection status  $s_{ij}$  (0: "retained" or 1: "discarded").  $T_{\text{RL}}$  is the upper bound of RL-steps. The observation  $O_\eta$  is obtained by  $O_\eta = [\mathbf{v}_i, \mathbf{v}_j, s_{ij, \eta}]$ . Note that the dimension of  $O_\eta$  only depends on the dimension of node attributes, which enables the applicability to the systems with varying numbers of entities. The policy network of RL-agent takes the observation  $O_\eta$  as input and decides the action at the current RL-step.

**Actions:** There are two possible actions for the RL-agent: "staying the same" (action 0) and "changing to the opposite" (action 1). At each RL-step, the RL-agent makes decision for each edge in the FC graph. The policy can be written as  $a = \pi(O)$ . We do not enforce any constraints on the selection of edges, i.e., there is no lower/upper bound on the number of selected edges. The actions of RL-agent may change the topology of the inferred graph  $\mathcal{G}'$  after each RL-step, which further influence the SGA-MG module.

**Rewards:** In general, the reward indicates how good the action taken by the RL-agent is with respect to the current situation. In our method, the reward is designed to indicate the performance of motion forecasting in various aspects. The acquisition of high rewards depends on the collaboration of all the modules in the framework.

The reward consists of three parts: regular reward  $R_{\text{reg}}$ , improvement reward  $R_{\text{imp}}$  and stimulation/punishment  $R_{\text{sti}}/R_{\text{pun}}$ . More specifically, the *regular reward* is the negative mean squared error of future predictions calculated by

$$R_{\text{reg}, \eta} = -\frac{1}{N} \sum_{i=1}^N \sum_{t'=t+1}^{t+T_f} \|\mathbf{x}_{t'}^i - \hat{\mathbf{x}}_{t', \eta}^i\|^2. \quad (5)$$

The *improvement reward* encourages the decrease of prediction error via applying a sign function to the error change between consecutive RL-steps, which is obtained by

$$R_{\text{imp}, \eta} = \text{sign}(R_{\text{reg}, \eta} - R_{\text{reg}, \eta-1}). \quad (6)$$

The reason of applying a sign function instead of directly using the raw improvement is to avoid reward vanishing when improvement becomes smaller towards convergence.

The *stimulation/punishment* is applied when there is a large improvement or deterioration in terms of a certain metric, which is given by

$$R_{\text{sti},\eta} = \Omega_s, \quad R_{\text{pun},\eta} = -\Omega_p, \quad (7)$$

where  $\Omega_s$  and  $\Omega_p$  are manually defined positive constants. These rewards depend on the metrics in specific domains.

Then the whole reward is calculated by

$$R_\eta = R_{\text{reg},\eta} + \beta_{\text{imp}} R_{\text{imp},\eta} + \beta_{\text{sti}} R_{\text{sti},\eta} \mathbb{I}(\text{sti}) + \beta_{\text{pun}} R_{\text{pun},\eta} \mathbb{I}(\text{pun}), \quad (8)$$

where  $\beta_{\text{imp}}$ ,  $\beta_{\text{sti}}$  and  $\beta_{\text{pun}}$  are hyperparameters and  $\mathbb{I}(\cdot)$  is an indicator function to indicate the occurrence of large improvement or deterioration.

#### 4.4. Soft Attention: Key Information Ranking

After finalizing the key information (edges) by the RL-HA module, the soft graph attention mechanism [53] is applied over the inferred graph  $\mathcal{G}'$  to further determine the relative significance of the selected key information at each time step. Here we take time step  $t$  as an example to illustrate the soft graph attention mechanism. Note that unlike the one-shot state embedding in GMP, here we construct a spatio-temporal graph to incorporate the state information with a stepwise embedding strategy.

In order to avoid confusion on notation with the RL-HA section, here we denote  $i$ -th node attribute at time  $t$  as

$$\bar{\mathbf{v}}_{i,t} = [\bar{\mathbf{v}}_{i,t}^{\text{self}}, \bar{\mathbf{v}}_{i,t}^{\text{social}}, \mathbf{v}_i^{\text{context}}], \quad (9)$$

where the context attribute  $\mathbf{v}_i^{\text{context}}$  is obtained in Section 4.2, the self-attribute  $\bar{\mathbf{v}}_{i,t}^{\text{self}}$  and the intermediate attribute  $\bar{\mathbf{v}}_{i,t}^{\text{neighbor}}$  are introduced in Section 4.5, and the social attribute  $\bar{\mathbf{v}}_{i,t}^{\text{social}}$  is calculated by the graph soft attention mechanism as follows

$$\bar{\alpha}_{ij}^t = \frac{\exp\left(\text{MLP}\left(\left[\bar{\mathbf{v}}_{i,t}^{\text{self}} \parallel \bar{\mathbf{v}}_{j,t}^{\text{neighbor}}\right]\right)\right)}{\sum_{k \in \mathcal{N}_i} \exp\left(\text{MLP}\left(\left[\bar{\mathbf{v}}_{i,t}^{\text{self}} \parallel \bar{\mathbf{v}}_{k,t}^{\text{neighbor}}\right]\right)\right)}, \quad (10)$$

$$\bar{\mathbf{v}}_{i,t}^{\text{social}} = \bar{f}_v \left( \sum_{j \in \mathcal{N}_i} \bar{\alpha}_{ij}^t \bar{\mathbf{v}}_{j,t}^{\text{neighbor}} \right),$$

where  $\bar{\alpha}_{ij}^t$  are learnable attention weights. We also use multi-head attention to stabilize training [53].

#### 4.5. Spatio-Temporal Motion Generator

The motion generator consists of two LSTM networks (E-LSTM/G-LSTM) with the soft graph attention in between. The E-LSTM takes in agent state information and outputs  $\bar{\mathbf{v}}_{i,t}^{\text{self}}$  at each time step, while the G-LSTM takes in the complete node attribute  $\bar{\mathbf{v}}_{i,t}$  and outputs the predicted change in state  $\Delta \hat{\mathbf{x}}_t^i$  at the current time  $t$ , which is used to calculate the state  $\hat{\mathbf{x}}_{t+1}^i$  with the system model (e.g.,

discrete-time linear dynamics). More specifically at time  $t$ ,

$$\text{Embedding: } \bar{\mathbf{v}}_{i,t}^{\text{self}} = \text{E-LSTM}^s(\mathbf{x}_t^i; \mathbf{h}_t^{s,i}), \quad (11)$$

$$\bar{\mathbf{v}}_{i,t}^{\text{neighbor}} = \text{E-LSTM}^n(\mathbf{x}_t^i; \mathbf{h}_t^{n,i}), \quad (12)$$

$$\text{Generation: } \Delta \hat{\mathbf{x}}_t^i = \text{G-LSTM}(\bar{\mathbf{v}}_{i,t}; \tilde{\mathbf{h}}_t^i), \quad (13)$$

$$\hat{\mathbf{x}}_{t+1}^i = f_{\text{system}}(\mathbf{x}_t^i, \Delta \hat{\mathbf{x}}_t^i) + \epsilon, \quad (14)$$

where  $\mathbf{h}_t^{s,i}$ ,  $\mathbf{h}_t^{n,i}$ , and  $\tilde{\mathbf{h}}_t^i$  are the hidden states of E-LSTM and G-LSTM respectively,  $\epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma)$  is a random noise to incorporate uncertainty, and  $f_{\text{system}}$  is the system model. Note that the whole generation process is divided into two stages: burn-in stage (from  $t - T_h + 1$  to  $t$ ) and prediction stage (from  $t + 1$  to  $t + T_f$ ). At burn-in stage, the true state is fed into E-LSTM while at prediction stage, the last prediction is fed instead.

If the topology of the inferred graph  $\mathcal{G}'$  is assumed to remain static over time, we can use one-shot generation to obtain the complete future trajectory. Otherwise, we can first generate the trajectory segment within a certain future horizon  $\tau < T_f$ , and push the predicted segment into the observations. This process can be iteratively propagated to generate the whole trajectory. We set  $\tau = 2$  in Figure 2.

#### 4.6. Framework Training Strategy

Since the RL-HA module is not differentiable, we cannot use an end-to-end fashion to train the whole model. Therefore, we propose an alternating training algorithm to improve RL-HA and the other modules separately. There are two training stages: pre-training stage and formal-training stage. In the *pre-training* stage, we pre-train the parameters of GMP module with an auto-encoder structure by unsupervised learning, where GMP serves as the encoder together with an auxiliary decoder. The decoder is discarded after the GMP is well trained. In addition, we pre-train the SGA-MG module with a fully connected topology to enable informative initial reward for the RL-HA module. The pre-training is necessary for the convergence of formal-training stage. In the *formal-training* stage, the pre-trained parameters of GMP are fixed to serve as a feature extractor, and we perform alternating optimization of RL-HA and SGA-MG modules: (a) train the RL-HA module with fixed parameters of GMP and SGA-MG with Double Deep Q-Learning (DDQN) method [18]; (b) finetune SGA-MG with fixed parameters of GMP and RL-HA using back-propagation methods. Please refer to Section 8 in the supplementary materials for the pseudo-code and implementation details.

### 5. RAIN for Human Skeleton Motions

In this section, we address the application of our method to capture long-term temporal dependency in skeleton based human motions. Instead of selecting important or relevant

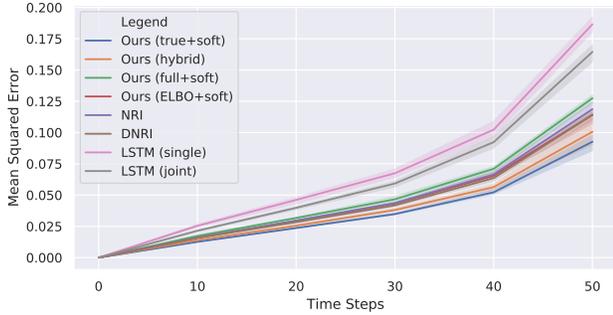


Figure 3. Mean squared error (MSE) in predicting future positions of mixed particles.

entities introduced in Section 4, here we propose to utilize the RL-HA module to discriminate key information over the whole history horizon. It can be either frame-wise selection or segment-wise selection. We demonstrate an illustrative case study of the latter based on a state-of-the-art model proposed in [59].

More specifically, the model in [59] is employed as the soft attention based motion generator in our framework. The proposed RL hard attention mechanism is built on top of it, and they work as a whole to generate future human motions. First, an encoding function is applied to extract contextual representation of each history frame. Then, the RL-HA module discriminates the key motion segments from history observations, and the soft attention mechanism further outputs relative significance of the selected segments for the current prediction. The model architecture and implementation details are elaborated in Section 9.

## 6. Experiments

In this paper, we validated the proposed hybrid attention approach on the datasets in different domains, including a synthetic physics motion dataset, a real-world driving dataset (nuScenes [4]) and a human motion dataset (Human3.6M [21]).

For the mixed particle simulation, since we have access to the ground truth of the particle charging state (i.e., positively/negatively-charged or uncharged), we can quantitatively and qualitatively evaluate the model performance in terms of both attention learning and mean square error (MSE) of particle position prediction. For the nuScenes driving dataset, we employ three widely used metrics: minimum average displacement error ( $\text{minADE}_{20}$ ), minimum final displacement error ( $\text{minFDE}_{20}$ ) and miss rate (MR) in trajectory prediction [26]. We report the results for vehicles and pedestrians separately. For the Human3.6M motion dataset, we adopt the standard metrics in the human pose estimation and motion forecasting literature [59], which is the Euclidean distance in Euler angle representation. We also provide ablative analysis on these datasets.

	Accuracy	Precision	Recall	F1-score
Corr. LSTM [23]	71.85±0.39	54.52±0.24	60.17±0.16	57.21±0.33
NRI [23]	85.47±0.98	78.57±0.59	64.71±0.66	70.97±0.51
DNRI [17]	87.54±1.66	82.48±0.91	67.23±0.78	74.08±0.83
Ours (ELBO+soft)	89.31±0.46	87.29±0.21	73.52±0.30	79.82±0.25
Ours (hybrid)	<b>93.52±1.25</b>	<b>92.34±0.87</b>	<b>80.55±0.92</b>	<b>86.04±0.74</b>
Supervised	97.48±0.23	95.24±0.09	89.55±0.14	92.31±0.11

Table 1. Evaluation (Mean±Std in %) of Relation Recognition (Mixed Particle System).

In all the experiments, we used a batch size of 32 and the Adam optimizer with an initial learning rate of 0.001 to train the models with a single NVIDIA Quadro RTX 6000 GPU. The additional experimental results, dataset details, baseline methods as well as implementation details are elaborated in supplementary materials.

### 6.1. Synthetic Simulation: Mixed Particle System

We applied our approach to a simulated physical system with mixed charged and uncharged particles, where the charged ones are randomly assigned positive or negative charges with a uniform distribution. The particles with the same charge repel and the ones with opposite charges attract according to the fundamental law of electricity. The uncharged particles perform uniform motion independently. The proposed hybrid-attention based prediction model is expected to learn the relations between particles and forecast their long-term motions.

The goal of RL-based hard attention mechanism is to infer whether there is a force between a pair of particles, which is essentially a binary classification task. The recognition of the force type (repel or attract) will be handled implicitly by the soft attention counterpart. The relation recognition results of different approaches are compared in Table 1, where the standard deviations are calculated based on three independent runs. It shows that the supervised learning method which directly trains a binary classifier with ground truth labels performs the best in terms of all the evaluation metrics. However, the accessibility of true relation labels is highly limited in real-world tasks. Among the approaches that do not require true labels during training, our method achieves the highest recognition accuracy, which improves by 6.8% over the strongest baseline DNRI and 4.7% over the ablative baseline *Ours (ELBO+soft)*. Based on the experiments, end-to-end training with ELBO tends to be more stable and converges faster than RL in general, but it did not achieve a better final performance. The accuracy alone cannot sufficiently prove the superiority due to the data imbalance of the two classes. Therefore, we also compared the precision, recall and F1-score. Our method achieves consistently better results on these metrics.

We further predicted the particle positions at the future 50 time steps given the historical observations of 30 time steps. The comparison of mean squared error (MSE) of

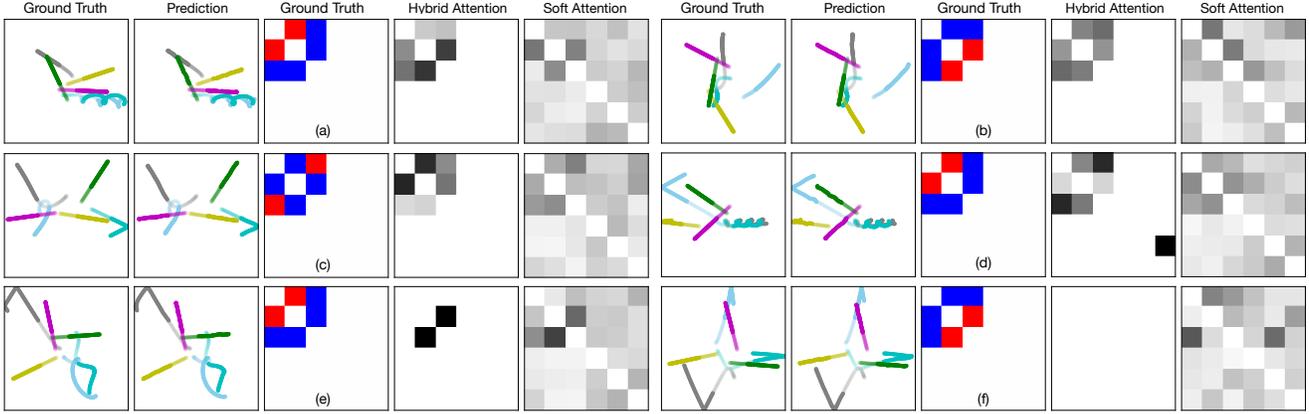


Figure 4. Visualization of particle trajectories and attention maps. The semi-transparent segments represent the observation of 30 time steps and the solid ones represent the prediction of future 50 time steps. In the ground truth relation map, red denotes “repel” and blue denotes “attract”. In the attention maps, darker color indicates larger weight. Note that we do not consider self-attention (i.e. zero diagonal).

Method	1.0s	2.0s	3.0s	4.0s	MR
NRI [23]	0.19/0.23	0.37/0.57	0.58/1.00	0.82/1.52	27.1
DNRI [17]	0.17/0.20	0.33/0.49	0.50/0.87	0.71/1.33	22.4
STGAT [20]	0.16/0.19	0.31/0.47	0.48/0.83	0.68/1.27	20.6
S-STGCNN [42]	0.18/0.21	0.34/0.52	0.53/0.92	0.75/1.41	24.7
Ours (full+soft)	0.07/0.10	0.17/0.41	0.40/0.82	0.63/1.28	21.0
Ours (ELBO+soft)	<b>0.04/0.07</b>	0.16/0.42	0.40/0.79	0.61/1.24	16.1
Ours (hybrid, static)	0.05/0.09	0.15/0.40	0.38/0.77	0.58/1.17	11.3
Ours (hybrid, dynamic)	0.06/0.09	<b>0.14/0.38</b>	<b>0.37/0.75</b>	<b>0.54/1.12</b>	<b>9.5</b>

Table 2. Comparison of  $\text{minADE}_{20} / \text{minFDE}_{20}$  (Meters) and Miss Rate@2.0m (MR, %) of Vehicle Trajectory Prediction.

different methods is provided in Figure 3. The solid lines and shaded areas represent the average value and standard deviation of three runs, respectively. *Ours (true+soft)* attains the smallest MSE consistently over the whole prediction horizon due to access to the true graph structures, which serves as a performance upper bound. *Ours (hybrid)* achieves a smaller MSE than *Ours (full+soft)* / *Ours (ELBO+soft)* and even approaches the performance of *Ours (true+soft)*, which indicates that the RL-based hard attention mechanism indeed helps on the relation recognition, which further improves prediction performance. *Ours (hybrid)* also outperforms NRI and DNRI where the interaction graph is encoded to a latent space.

We also visualize the predicted trajectories and attention maps of typical cases in Figure 4 with descriptions in the caption. The results show that the RL agent can figure out truly relevant agents with a high accuracy and the soft attention counterpart further determines the relative significance of the selected entities. However, applying soft attention alone on fully connected graphs sometimes assigns trivial weights or even large weights to irrelevant agents as shown in the “Soft Attention” columns, which weakens interpretability and performance of the model. We address two interesting cases in (e) and (f), where the RL agent only selects a subset of interacting agents or even an empty set, but the system still generates very good prediction. The rea-

Method	1.0s	2.0s	3.0s	4.0s	MR
NRI [23]	0.10/0.13	0.22/0.32	0.35/0.58	0.48/0.82	28.5
DNRI [17]	0.09/0.12	0.20/0.31	0.32/0.53	0.44/0.78	25.6
STGAT [20]	0.08/0.11	0.18/0.27	0.29/0.48	0.40/0.71	17.8
S-STGCNN [42]	0.12/0.15	0.23/0.37	0.35/0.61	0.48/0.88	27.9
Ours (full+soft)	0.07/0.11	0.15/0.22	0.24/0.40	0.32/0.56	13.2
Ours (ELBO+soft)	<b>0.04/0.06</b>	0.13/0.20	0.21/0.39	0.30/0.56	12.9
Ours (hybrid, static)	0.06/0.08	0.12/0.20	0.19/0.36	0.29/0.54	11.7
Ours (hybrid, dynamic)	0.05/0.08	<b>0.11/0.18</b>	<b>0.17/0.34</b>	<b>0.26/0.51</b>	<b>9.8</b>

Table 3. Comparison of  $\text{minADE}_{20} / \text{minFDE}_{20}$  (Meters) and Miss Rate@1.0m (MR, %) of Pedestrian Trajectory Prediction.

son is that the distance between the interacting particles is too large, which leads to the ignorance of the weak forces by the RL agent. Particularly in (f), all the particles perform nearly uniform motions which results in an empty hybrid attention map. This demonstrates the capability of distinguishing key information of the hard attention mechanism.

## 6.2. nuScenes Dataset: Traffic Scenarios

We validated our RAIN framework on the nuScenes dataset, which handles long-term prediction of heterogeneous traffic participants (i.e., vehicles and pedestrians). We predicted the future 4.0s (20 frames) given the history observations of 2.0s (10 frames). The comparison of quantitative results is shown in Table 2 (vehicles) and Table 3 (pedestrians). Note that we simultaneously included both types of agents during training, while reported the test results separately as done in [46].

All the learning-based baseline methods consider the relation modeling via certain techniques. It is shown that *Ours (hybrid, dynamic)* achieves the smallest error and lowest miss rate in both vehicle and pedestrian prediction. We also present extensive ablation results. *Ours (full+soft)* achieves lower prediction error than STGAT with a similar model structure, especially in vehicle prediction. The reason is that the vehicle kinematics model is incorporated into the motion generator to propagate the vehicle state, which can

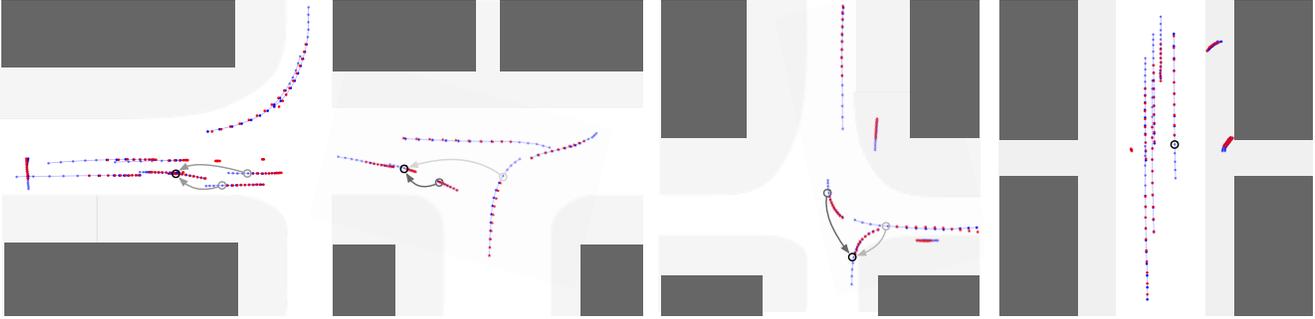


Figure 5. The visualization of testing cases in the nuScenes dataset. The light blue dots are history observations, dark blue dots are ground truth, and red dots are predictions with the minimum ADE. The black circles indicate the target agent and gray arrows indicate hybrid attention. The targets only attend to the agents with arrows selected by hard attention and darker colors imply larger weights. In the last case, there is no arrow, which implies that the model infers that the target is not influenced by any agent in the scene at the current frame.

Motion milliseconds	Walking					Eating					Smoking					Discussion				
	80	160	320	400	1k	80	160	320	400	1k	80	160	320	400	1k	80	160	320	400	1k
Res-sup. [40]	0.27	0.46	0.67	0.75	1.03	0.23	0.37	0.59	0.73	1.08	0.32	0.59	1.01	1.10	1.50	0.30	0.67	0.98	1.06	1.69
CSM [27]	0.33	0.54	0.68	0.73	0.92	0.22	0.36	0.58	0.71	1.24	0.26	0.49	0.96	0.92	1.62	0.32	0.67	0.94	1.01	1.86
Traj-GCN [34]	0.18	0.32	0.49	0.56	0.67	0.17	0.31	0.52	0.62	1.12	0.22	0.41	0.84	0.79	1.57	0.20	<b>0.51</b>	0.79	0.86	1.70
DMGNN [31]	0.18	0.31	0.49	0.58	0.75	0.17	0.30	0.49	0.59	1.14	0.21	0.39	0.81	0.77	1.52	0.26	0.65	0.92	0.99	<b>1.45</b>
LTD-10-10 [39]	0.18	0.31	0.49	0.56	0.64	<b>0.16</b>	0.29	0.50	0.62	1.10	0.22	0.41	0.86	0.80	1.58	0.20	<b>0.51</b>	<b>0.77</b>	<b>0.85</b>	1.75
HisReplself [59]	0.18	0.30	0.46	0.51	0.64	<b>0.16</b>	0.29	0.49	0.60	1.10	0.22	0.42	0.86	0.80	1.58	0.20	0.52	0.78	0.87	1.63
Ours (hybrid)	<b>0.17</b>	<b>0.28</b>	<b>0.43</b>	<b>0.49</b>	<b>0.60</b>	<b>0.16</b>	<b>0.26</b>	<b>0.45</b>	<b>0.56</b>	<b>1.01</b>	<b>0.20</b>	<b>0.38</b>	<b>0.80</b>	<b>0.74</b>	<b>1.48</b>	<b>0.19</b>	0.52	0.82	0.91	1.64

Table 4. Comparison of mean angle errors (MAE) of different methods for both short-term ( $\leq 400$  milliseconds) and long-term prediction (1k milliseconds) on four representative actions in the Human3.6M dataset.

guarantee the feasibility of generated trajectories. *Ours (hybrid, static)* and *Ours (ELBO+soft)* outperforms *Ours (full+soft)*, which implies the effectiveness of hard attention on selecting important agents. In particular, *Ours (hybrid, static)* achieves smaller prediction errors than *Ours (ELBO+soft)*, which implies the advantage of RL based hard attention over the ELBO based method. *Ours (hybrid, dynamic)* further reduces the error owing to the dynamic recognition of key elements, where the margin of improvement becomes larger as the prediction horizon increases.

The visualization of agent trajectories and hybrid attention weights in several typical test scenarios is provided in Figure 5. It is shown that our method can learn reasonable attention weights to exploit the key information, and generate plausible and accurate prediction hypotheses.

### 6.3. Human3.6M Dataset: Human Motions

We adopted the same experimental setting as the baseline approaches for fair comparison, which is to forecast the future 25 frames based on the past 50 frames. The model was trained to only predict 10 frames during the training phase, while the prediction was recursively applied by using predictions as new observations during the testing phase. The forecasting results are shown in Table 4, where we only include four representative actions due to space limit. The complete table can be found in Table 5 in the supplementary materials. Among the baselines, HisReplself [59]

yields the previous state-of-the-art performance, which also serves as an ablation model with only soft attention on the motion history. In general, the results show that *Ours (hybrid)* achieves the best performance both in average and for most actions in terms of both short-term and long-term forecasting accuracy. In particular, *Ours (hybrid)* outperforms HisReplself in average, indicating the additional benefits brought by the hard attention mechanism.

## 7. Conclusion

In this paper, we present a generic motion forecasting framework with key information selection and ranking procedures. The former is realized by a reinforcement learning based hard attention mechanism for the purpose of discriminating relevant information for the prediction from complete spatio-temporal observations. The latter is fulfilled by a soft attention mechanism for the purpose of determining relative importance of the selected key elements. These dual procedures can be applied recurrently to dynamically adjust the focus of the model as the situation evolves over time. A double-stage training pipeline with an alternating training strategy is employed to train different parts of the model, which proves to be effective and stable in the experiments. The general framework is instantiated and applied to multi-agent trajectory prediction and human motion forecasting, which achieves state-of-the-art performance in terms of a wide range of evaluation metrics in different domains.

## References

- [1] Alexandre Alahi, Kratharth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016. 2
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. 3
- [3] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, pages 4502–4510, 2016. 2
- [4] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019. 6, 16
- [5] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In *Conference on Robot Learning*, pages 86–99, 2020. 2
- [6] Sangwoo Cho, Muhammad Maqbool, Fei Liu, and Hassan Foroosh. Self-attention network for skeleton-based human action recognition. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 635–644, 2020. 3
- [7] Chiho Choi, Joon Hee Choi, Jiachen Li, and Srikanth Malla. Shared cross-modal trajectory prediction for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2021. 2
- [8] Chiho Choi and Behzad Dariush. Looking to relations for future trajectory forecast. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 921–930, 2019. 2
- [9] Nachiket Deo and Mohan M Trivedi. Convolutional social pooling for vehicle trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1468–1476, 2018. 2
- [10] Chenyou Fan, Yuze Zhang, Yi Pan, Xiaoyue Li, Chi Zhang, Rong Yuan, Di Wu, Wensheng Wang, Jian Pei, and Heng Huang. Multi-horizon time series forecasting with temporal attention learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2527–2535, 2019. 3
- [11] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4346–4354, 2015. 2
- [12] Luca Franceschi, Mathias Niepert, Massimiliano Pontil, and Xiao He. Learning discrete structures for graph neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 1972–1982, 2019. 2
- [13] Haiming Gang, Harshayu Girase, Srikanth Malla, Jiachen Li, Akira Kanehara, Karttikeya Mangalam, and Chiho Choi. Loki: Long term and key intentions for trajectory prediction. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 2
- [14] Hongyang Gao and Shuiwang Ji. Graph representation learning via hard and channel-wise attention networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 741–749, 2019. 2
- [15] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11525–11533, 2020. 2
- [16] Anand Gopalakrishnan, Ankur Mali, Dan Kifer, Lee Giles, and Alexander G Ororbia. A neural temporal model for human motion prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12116–12125, 2019. 2
- [17] Colin Graber and Alexander G Schwing. Dynamic neural relational inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8513–8522, 2020. 6, 7, 17
- [18] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 2094–2100. AAAI Press, 2016. 5, 14
- [19] Yedid Hoshen. Vain: Attentional multi-agent predictive modeling. In *Advances in Neural Information Processing Systems*, pages 2701–2711, 2017. 2
- [20] Yingfan Huang, Huikun Bi, Zhaoxin Li, Tianlu Mao, and Zhaoqi Wang. Stgat: Modeling spatial-temporal interactions for human trajectory prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6272–6281, 2019. 2, 3, 7
- [21] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014. 6, 16
- [22] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5308–5317, 2016. 2
- [23] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In *International Conference on Machine Learning*, pages 2688–2697, 2018. 2, 6, 7, 16, 17
- [24] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezaeifighi, and Silvio Savarese. Socialbigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks. In *Advances in Neural Information Processing Systems*, pages 137–146, 2019. 2, 3
- [25] Jannik Kossen, Karl Stelzner, Marcel Hussing, Claas Voelcker, and Kristian Kersting. Structured object-aware physics

- prediction for video modeling and planning. In *International Conference on Learning Representations*, 2020. 2
- [26] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 336–345, 2017. 6
- [27] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5226–5234, 2018. 2, 8, 14
- [28] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6150–6156. IEEE, 2019. 2, 3
- [29] Jiachen Li, Hengbo Ma, Zhihao Zhang, Jinning Li, and Masayoshi Tomizuka. Spatio-temporal graph dual-attention network for multi-agent prediction and tracking. *IEEE Transactions on Intelligent Transportation Systems*, 2021. 2, 3
- [30] Jiachen Li, Fan Yang, Masayoshi Tomizuka, and Chiho Choi. Evolvegraph: Multi-agent trajectory prediction with dynamic relational reasoning. In *Advances in neural information processing systems*, 2020. 2
- [31] Maosen Li, Siheng Chen, Yangheng Zhao, Ya Zhang, Yanfeng Wang, and Qi Tian. Dynamic multiscale graph neural networks for 3d skeleton based human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 214–223, 2020. 8, 14
- [32] Yunzhu Li, Antonio Torralba, Anima Anandkumar, Dieter Fox, and Animesh Garg. Causal discovery in physical systems from videos. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [33] Nian Liu, Junwei Han, and Ming-Hsuan Yang. Picanet: Learning pixel-wise contextual attention for saliency detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3089–3098, 2018. 3
- [34] Zhenguang Liu, Shuang Wu, Shuyuan Jin, Qi Liu, Shijian Lu, Roger Zimmermann, and Li Cheng. Towards natural and accurate future motion prediction of humans and animals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10004–10012, 2019. 8, 14
- [35] Hengbo Ma, Yaofeng Sun, Jiachen Li, and Masayoshi Tomizuka. Multi-agent driving behavior prediction across different scenarios with self-supervised domain knowledge. In *2021 International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2021. 2
- [36] Hengbo Ma, Yaofeng Sun, Jiachen Li, Masayoshi Tomizuka, and Chiho Choi. Continual multi-agent interaction behavior prediction with conditional generative memory. *IEEE Robotics and Automation Letters*, 2021. 2
- [37] Xiaobai Ma, Jiachen Li, Mykel J Kochenderfer, David Isele, and Kikuo Fujimura. Reinforcement learning for autonomous driving with latent state inference and spatial-temporal relationships. In *2021 International Conference on Robotics and Automation (ICRA)*. IEEE, 2021. 2
- [38] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017. 2
- [39] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9489–9497, 2019. 2, 8, 14
- [40] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2891–2900, 2017. 8, 14
- [41] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015. 2
- [42] Abdullh Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14424–14432, 2020. 7
- [43] Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, 2016. 3
- [44] Tung Phan-Minh, Elena Corina Grigore, Freddy A Boulton, Oscar Beijbom, and Eric M Wolff. Covernet: Multimodal behavior prediction using trajectory sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14074–14083, 2020. 2
- [45] Yao Qin, Dongjin Song, Haifeng Cheng, Wei Cheng, Guofei Jiang, and Garrison W Cottrell. A dual-stage attention-based recurrent neural network for time series prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 2627–2633, 2017. 3
- [46] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control. In *Proceedings of Europe Conference on Computer Vision (ECCV)*, 2020. 7
- [47] Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, Sen Wang, and Chengqi Zhang. Reinforced self-attention network: a hybrid of hard and soft attention for sequence modeling. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4345–4352, 2018. 2, 3
- [48] Chenyang Si, Wentao Chen, Wei Wang, Liang Wang, and Tieniu Tan. An attention enhanced graph convolutional lstm network for skeleton-based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1227–1236, 2019. 3
- [49] Sijie Song, Culing Lan, Junliang Xing, Wenjun Zeng, and Jiaying Liu. Spatio-temporal attention-based lstm networks for 3d action recognition and detection. *IEEE Transactions on image processing*, 27(7):3459–3471, 2018. 3

- [50] Charlie Tang and Russ R Salakhutdinov. Multiple futures prediction. In *Advances in Neural Information Processing Systems*, pages 15424–15434, 2019. 2
- [51] Yongyi Tang, Lin Ma, Wei Liu, and Wei-Shi Zheng. Long-term human motion prediction by modeling motion context and enhancing motion dynamic. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 935–941, 2018. 2
- [52] Sjoerd van Steenkiste, Michael Chang, Klaus Greff, and Jürgen Schmidhuber. Relational neural expectation maximization: Unsupervised discovery of objects and their interactions. In *International Conference on Learning Representations*, 2018. 2
- [53] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018. 5
- [54] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *2018 IEEE international Conference on Robotics and Automation (ICRA)*, pages 1–7. IEEE, 2018. 2, 3
- [55] Boyu Wang, Lihan Huang, and Minh Hoai. Active vision for early recognition of human actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1081–1091, 2020. 2
- [56] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2017. 3
- [57] Yequan Wang, Minlie Huang, Xiaoyan Zhu, and Li Zhao. Attention-based lstm for aspect-level sentiment classification. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 606–615, 2016. 3
- [58] Nicholas Watters, Daniel Zoran, Theophane Weber, Peter Battaglia, Razvan Pascanu, and Andrea Tacchetti. Visual interaction networks: Learning a physics simulator from video. In *Advances in neural information processing systems*, pages 4539–4547, 2017. 2
- [59] Mao Wei, Liu Miaomiao, and Salzemann Mathieu. History repeats itself: Human motion prediction via motion attention. In *ECCV*, 2020. 2, 3, 6, 8, 12, 13, 14, 16, 17
- [60] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992. 2
- [61] Pengxiang Wu, Siheng Chen, and Dimitris N Metaxas. Motionnet: Joint perception and motion prediction for autonomous driving based on bird’s eye view maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11385–11395, 2020. 2
- [62] Wenhao Wu, Dongliang He, Xiao Tan, Shifeng Chen, and Shilei Wen. Multi-agent reinforcement learning based frame sampling for effective untrimmed video recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6222–6231, 2019. 2
- [63] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015. 1, 3
- [64] Ting Zhao and Xiangqian Wu. Pyramid feature attention network for saliency detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3085–3094, 2019. 3