

# Neural Relational Inference with Efficient Message Passing Mechanisms

Siyuan Chen, Jiahai Wang\*, Guoqing Li

School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China  
 {chensy47, ligq7}@mail2.sysu.edu.cn, wangjiah@mail.sysu.edu.cn

## Abstract

Many complex processes can be viewed as dynamical systems of interacting agents. In many cases, only the state sequences of individual agents are observed, while the interacting relations and the dynamical rules are unknown. The neural relational inference (NRI) model adopts graph neural networks that pass messages over a latent graph to jointly learn the relations and the dynamics based on the observed data. However, NRI infers the relations independently and suffers from error accumulation in multi-step prediction at dynamics learning procedure. Besides, relation reconstruction without prior knowledge becomes more difficult in more complex systems. This paper introduces efficient message passing mechanisms to the graph neural networks with structural prior knowledge to address these problems. A relation interaction mechanism is proposed to capture the coexistence of all relations, and a spatio-temporal message passing mechanism is proposed to use historical information to alleviate error accumulation. Additionally, the structural prior knowledge, symmetry as a special case, is introduced for better relation prediction in more complex systems. The experimental results on simulated physics systems show that the proposed method outperforms existing state-of-the-art methods.

## Introduction

Many complex processes in natural and social areas including multi-agent systems (Yang et al. 2018; Li et al. 2020), swarm systems (Oliveira et al. 2020), physical systems (Ha and Jeong 2020; Bapst et al. 2020) and social systems (Almaatouq et al. 2020; Zhang et al. 2020) can be viewed as dynamical systems of interacting agents. Revealing the underlying interactions and dynamics can help us understand, predict, and control the behavior of these systems. However, in many cases, only the state sequences of individual agents are observed, while the interacting relations and the dynamical rules are unknown.

A lot of works (Hoshen 2017; van Steenkiste et al. 2018; Watters et al. 2017) use implicit interaction models to learn the dynamics. These models can be regarded as graph neural networks (GNNs) over a fully-connected graph, and the implicit interactions are modeled via message passing operations (Watters et al. 2017) or attention mechanisms (Hoshen

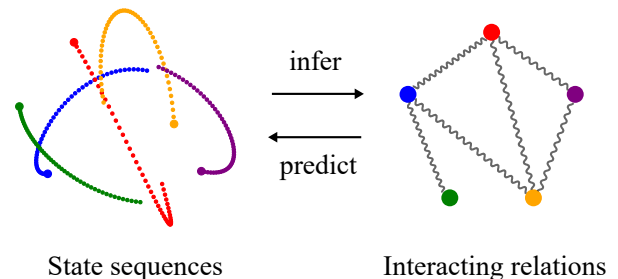


Figure 1: A dynamical system consisting of 5 particles that are linked by invisible springs. The interacting relations and future states are to be predicted based on the observed state sequences.

2017). Compared with modeling implicit interactions, modeling the explicit interactions offers a more interpretable way to understand the dynamical systems. A motivating example is shown in Fig. 1 (Kipf et al. 2018), where a dynamical system consists of 5 particles linked by invisible springs. It is of interest to infer the relations among the particles and predict their future states. Kipf et al. (2018) propose the neural relational inference (NRI) model, a variational auto-encoder (VAE) (Kingma and Welling 2014), to jointly learn explicit interactions and dynamical rules in an unsupervised manner.

Currently, there are three main limitations of NRI. First, the interacting relations are inferred independently, while the coexistence of these relations is not considered. Alet et al. (2019) tackle this problem by taking all relations as a whole and iteratively improving the prediction through modular meta-learning. However, this method is computationally costly and it is limited to small-scale systems. Second, NRI predicts multiple steps into the future to emphasize the effect of the interactions, which leads to error accumulation and prevents the model from reconstructing the dynamics accurately. Third, as the scale and the complexity of systems increase, it is difficult to infer the interactions solely based on the observed data, while incorporating some structural prior knowledge can help better explore the structural space and promote the precision of relation recovery.

To address the problems above, this paper introduces efficient message passing mechanisms with structural prior knowledge for neural relational inference. For relation re-

\*Corresponding author

construction, a relation interaction mechanism is introduced to capture the dependencies among different relations. For dynamics reconstruction, a spatio-temporal message passing mechanism is introduced to utilize historical information to alleviate the error accumulation of multi-step prediction. Additionally, the prior knowledge about the relations is incorporated as a regularization term in the loss function to impose soft constraints, and as a special case, the symmetry of relations is taken into consideration.

The contributions of this work are summarized as follows.

- Efficient message passing mechanisms are introduced to make a joint prediction of relations and alleviate the error accumulation of multi-step state prediction.
- The prior knowledge about relations, symmetry as a special case, is introduced to better reconstruct relations on more complex dynamical systems.
- Extensive experiments on physics simulation datasets are conducted. The results show the superiority of our method by comparing it with the state-of-the-art methods.

## Background

### Neural Relational Inference (NRI)

NRI (Kipf et al. 2018) is an unsupervised model that infers interacting structure from observational data and learns system dynamics. NRI takes the form of VAE, where the encoder infers the interacting relations, and the decoder predicts the future states of individual agents.

Specifically, the encoder adopts GNNs with multiple rounds of message passing, and infers the distribution of potential interactions based on the input state,

$$\mathbf{h} = f_{\text{enc}}(\mathbf{x}), \quad (1)$$

$$q_{\phi}(\mathbf{z}|\mathbf{x}) = \text{softmax}(\mathbf{h}), \quad (2)$$

where  $\mathbf{x} = (\mathbf{x}_1^{1:T}, \dots, \mathbf{x}_N^{1:T})$  is the observed trajectories of  $N$  objects in the system in  $T$  time steps,  $f_{\text{enc}}$  is a GNN acting on the fully-connected graph, and  $q_{\phi}(\mathbf{z}|\mathbf{x})$  is the factorized distribution of edge type  $\mathbf{z}$ .

Since directly sampling edges from  $q_{\phi}(\mathbf{z}|\mathbf{x})$ , a discrete distribution, is a non-differentiable process, the back-propagation cannot be used. To solve this problem, NRI uses the Gumbel-Softmax trick (Maddison, Mnih, and Teh 2017), which simulates a differentiable sampling process from a discrete distribution using a continuous function, i.e.,

$$\mathbf{z}_{ij} = \text{softmax}((\mathbf{h}_{ij} + \mathbf{g})/\tau), \quad (3)$$

where  $\mathbf{z}_{ij}$  represents the edge type between the nodes  $v_i$  and  $v_j$ ,  $\mathbf{g} \in \mathbb{R}^K$  is a vector of i.i.d. samples from a Gumbel(0, 1) distribution and the temperature  $\tau$  is a parameter controlling the ‘‘smoothness’’ of the samples.

According to the inferred relations and past state sequences, the decoder uses another GNN that models the effect of interactions to predict the future state,

$$p_{\theta}(\mathbf{x}|\mathbf{z}) = \prod_{t=1}^T p_{\theta}(\mathbf{x}^{t+1}|\mathbf{x}^{1:t}, \mathbf{z}), \quad (4)$$

where  $p_{\theta}(\mathbf{x}^{t+1}|\mathbf{x}^{1:t}, \mathbf{z})$  is the conditional likelihood of  $\mathbf{x}^{t+1}$ .

As a variational auto-encoder model, NRI is trained to maximize the evidence lower bound,

$$\mathcal{L} = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - \text{KL}[q_{\phi}(\mathbf{z}|\mathbf{x})||p_{\theta}(\mathbf{z})], \quad (5)$$

where the prior  $p_{\theta}(\mathbf{z}) = \prod_{i \neq j} p_{\theta}(\mathbf{z}_{ij})$  is a factorised uniform distribution over edge types. In the right-hand side of Eq. (5), the first term is the expected reconstruction error, while the second term encourages the approximate posterior  $q_{\phi}(\mathbf{z}|\mathbf{x})$  to approach the prior  $p_{\theta}(\mathbf{z})$ .

### Message Passing Mechanisms in GNNs

GNNs are a widely used class of neural networks that operates on graph structured data by message passing mechanisms (Gilmer et al. 2017). For a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with vertices  $v \in \mathcal{V}$  and edges  $e = (v, v') \in \mathcal{E}$ , the node-to-edge ( $n \rightarrow e$ ) and edge-to-node ( $e \rightarrow n$ ) message passing operations are defined as follows,

$$v \rightarrow e: \quad \mathbf{h}_{(i,j)}^l = f_e^l([\mathbf{h}_i^l, \mathbf{h}_j^l, \mathbf{r}_{(i,j)}]), \quad (6)$$

$$e \rightarrow v: \quad \mathbf{h}_j^{l+1} = f_v^l([\sum_{i \in \mathcal{N}_j} \mathbf{h}_{(i,j)}^l, \mathbf{x}_j]), \quad (7)$$

where  $\mathbf{h}_i^l$  is the embedding of node  $v_i$  in the  $l$ -th layer of the GNNs,  $\mathbf{r}_{ij}$  is the feature of edge  $e_{ij}$  (e.g. edge type), and  $\mathbf{h}_{ij}^l$  is the embedding of edge  $e_{ij}$  in the  $l$ -th layer of the GNNs.  $\mathcal{N}_j$  denotes the set of indices of neighboring nodes with an incoming edge connected to vertex  $v_j$ , and  $[\cdot, \cdot]$  denotes the concatenation of vectors.  $f_e$  and  $f_v$  are node-specific and edge-specific neural networks, such as multilayer perceptrons (MLP), respectively. Node embeddings are mapped to edge embeddings in Eq. (6) and vice versa in Eq. (7).

## Method

The structure of our method is shown in Fig. 2. Our method follows the framework of VAE in NRI. In the encoder, a relation interaction mechanism is used to capture the dependencies among the latent edges for a joint relation prediction. In the decoder, a spatio-temporal message passing mechanism is used to incorporate historical information for alleviating the error accumulation in multi-step prediction. As both mechanisms mentioned above can be regarded as Message Passing Mechanisms in GNNs for Neural Relational Inference, our method is named as NRI-MPM. Additionally, the structural prior knowledge, symmetry as a special case, is incorporated as a regularization term in the loss function to improve relation prediction in more complex systems.

### Relation Interaction Mechanism

The encoder is aimed at inferring the edge types  $\mathbf{z}_{ij}$  based on the observed state sequences  $\mathbf{x} = (\mathbf{x}^1, \dots, \mathbf{x}^T)$ . From the perspective of message passing, the encoder defines a node-to-edge message passing operation at a high level. As shown in Eqs. (8)-(11), NRI first maps the observed state  $\mathbf{x}_j$  to a latent vector  $\mathbf{h}_j^1$ , and then applies two rounds of node-to-edge and one round of edge-to-node message passing alternately to obtain the edge embeddings  $\mathbf{h}_{(i,j)}^2$  that integrate both local and global information. Then,  $\mathbf{h}_{(i,j)}^2$  are used to predict

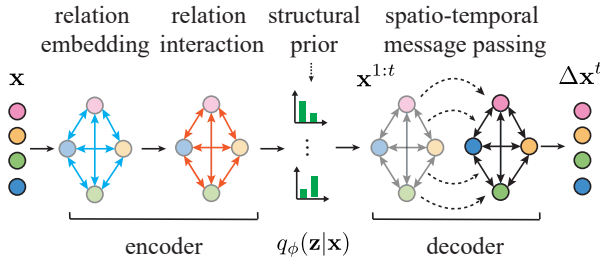


Figure 2: Overview of NRI-MPM. The encoder  $q_\phi(\mathbf{z}|\mathbf{x})$  uses the state sequences  $\mathbf{x}$  to generate relation embeddings, and applies a relation interaction mechanism to jointly predict the relations. The structural prior knowledge of symmetry is imposed as a soft constraint for relation prediction. The decoder takes the predicted relations  $\mathbf{z}$  and the historical state sequences  $\mathbf{x}^{1:t}$  to predict the change in state  $\Delta\mathbf{x}^t$ .

the pairwise interacting types independently.

$$\mathbf{h}_j^1 = f_{\text{emb}}(\mathbf{x}_j), \quad (8)$$

$$v \rightarrow e: \quad \mathbf{h}_{(i,j)}^1 = f_e^1([\mathbf{h}_i^1, \mathbf{h}_j^1]), \quad (9)$$

$$e \rightarrow v: \quad \mathbf{h}_j^2 = f_v^1\left(\sum_{i \neq j} \mathbf{h}_{(i,j)}^1\right), \quad (10)$$

$$v \rightarrow e: \quad \mathbf{h}_{(i,j)}^2 = f_e^2([\mathbf{h}_i^2, \mathbf{h}_j^2]). \quad (11)$$

However, the relations are generally dependent on each other since they jointly affect the future states of individual agents. Within the original formulations of GNNs, Eqs. (8)-(11) cannot effectively modeled the dependencies among edges. Typical GNNs are designed to learn node embeddings, while the edge embeddings are treated as a transient part of the computation. To capture the coexistence of all relations, this paper introduces an edge-to-edge ( $e \rightarrow e$ ) message passing operation that directly passes messages among edges, named the relation interaction mechanism,

$$e \rightarrow e: \{ \mathbf{e}_{(i,j)} : (v_i, v_j) \in \mathcal{E} \} = g_e \left( \{ \mathbf{h}_{(i,j)}^2 : (v_i, v_j) \in \mathcal{E} \} \right). \quad (12)$$

Ideally, this operation includes modeling the pairwise dependencies among all edges, which is computationally costly as its time complexity is  $O(|\mathcal{E}|^2)$ . Alternatively, as shown in Fig. 3, our method decomposes this operation into two sub-operations, intra-edge interaction and inter-edge interaction operations, for modeling the interactions among incoming edges of the same node and those among the incoming edges of different nodes, respectively.

The intra-edge interaction operation is defined as follows,

$$e \rightarrow e: \{ \mathbf{e}_{(i,j)}^1 : i \in \mathcal{N}_j \} = g_e^{\text{intra}} \left( \{ \mathbf{h}_{(i,j)}^2 : i \in \mathcal{N}_j \} \right). \quad (13)$$

From the above definition,  $g_e^{\text{intra}}$  is required to be permutation equivariant (Lee et al. 2019) to preserve the correspondences between the input and output edge embeddings. Formally, let  $\mathcal{I}_j = \{i_k\}_{k=1}^{|\mathcal{N}_j|}$  be the ascending sequence of  $\mathcal{N}_j$  and  $\mathcal{S}_j$  be the set of all permutations on  $\mathcal{I}_j$ . For any permutation  $\pi \in \mathcal{S}_j$ ,  $\pi(\mathbf{h}_{(\mathcal{I}_j, j)}^2) = \mathbf{h}_{(\pi(\mathcal{I}_j), j)}^2$ .  $g_e^{\text{intra}}$  is said to be permutation equivariant if  $\pi(\mathbf{e}_{(\mathcal{I}_j, j)}^1) = g_e^{\text{intra}} \circ \pi(\mathbf{h}_{(\mathcal{I}_j, j)}^2)$ ,

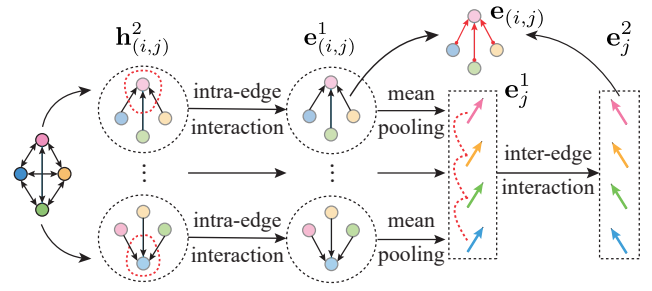


Figure 3: Relation interaction mechanism. Given the edge embeddings  $\mathbf{h}_{(i,j)}^2$ , the intra-edge and inter-edge interaction operations are used to model the interactions among incoming edges of the same node and those among the incoming edges of different nodes, respectively. The resulting embeddings  $\mathbf{e}_{(i,j)}^1$  and  $\mathbf{e}_j^2$  are concatenated to obtain the final edge representations  $\mathbf{e}_{(i,j)}$ .

where  $\circ$  denotes the composition of functions. Although permutation invariant operators such as sum and max aggregators are widely used in GNNs (Xu et al. 2019), the design for permutation equivariant operators is less explored. Inspired by Hamilton et al. (2017), this paper treats a set as an unordered sequence and defines  $g_e^{\text{intra}}$  as a sequence model over it, appending by an inverse permutation to restore the order, i.e.,

$$e \rightarrow e: \quad \mathbf{e}_{(\mathcal{I}_j, j)}^1 = \pi^{-1} \circ S \circ \pi \left( \mathbf{h}_{(\mathcal{I}_j, j)}^2 \right), \quad (14)$$

where  $S$  is a sequence model such as recurrent neural networks (RNNs) and convolutional neural networks (CNNs). When  $S$  is implemented by self-attention,  $g_e^{\text{intra}}$  is permutation equivariant, which may not hold when  $S$  is an RNN. In this case,  $g_e^{\text{intra}}$  is expected to be approximately permutation equivariant with a well trained  $S$ .

Similarly, one can define an inter-edge interaction operation. The difference is that this operation treats all incoming edges of a node as a whole. For simplicity, this paper applies mean-pooling to the edge embeddings to get an overall representation. Formally, the inter-edge interaction operation is defined to be the composition of the following steps,

$$e \rightarrow v: \quad \mathbf{e}_j^1 = \text{Mean} \left( \mathbf{e}_{(1:N, j)}^1 \right), \quad (15)$$

$$v \rightarrow v: \quad \{ \mathbf{e}_j^2 \}_{j=1}^N = g_e^{\text{inter}} \left( \{ \mathbf{e}_j^1 \}_{j=1}^N \right), \quad (16)$$

where  $g_e^{\text{inter}}$  is a node-to-node ( $v \rightarrow v$ ) operation that passes messages among nodes and takes a similar form of  $g_e^{\text{intra}}$ . To analyze the time complexity of relation interaction, this paper assumes that all sequence models involved are RNNs. Since  $g_e^{\text{intra}}$  and Mean can be applied to each node in parallel and the time complexity of Eqs. (14)-(16) are all  $O(N)$ , the overall time complexity of relation interaction is  $O(N)$ , which is much more effective than the pairwise interaction.

Finally, an MLP is used to unify the results of the two operations, and the predicted edge distribution is defined as

$$\mathbf{e}_{(i,j)} = \text{MLP}([\mathbf{e}_{(i,j)}^1, \mathbf{e}_j^2]), \quad (17)$$

$$q_\phi(\mathbf{z}_{ij}|\mathbf{x}) = \text{softmax}(\mathbf{e}_{(i,j)}). \quad (18)$$

As in Eq. (3),  $\mathbf{z}_{ij}$  is sampled via the Gumbel-Softmax trick to allow back-propagation.

### Spatio-Temporal Message Passing Mechanism

The decoder is aimed at predicting the future state  $\mathbf{x}^{t+1}$  using the inferred relations  $\mathbf{z}$  and the historical states  $\mathbf{x}^{1:t}$ . Since the interactions can have a small effect on short-term dynamics, NRI predicts multiple steps into the future to avoid a degenerate decoder that ignores the effect of interactions. In multi-step prediction, the predicted value  $\mu_j^t$  is replaced by the ground truth state  $\mathbf{x}_j^t$ , i.e.,

$$\mu_j^{t+1} = f_{\text{dec}}(\mathbf{x}_j^t), \quad (19)$$

$$\mu_j^{t+m} = f_{\text{dec}}(\mu_j^{t+m-1}), \quad m = 2, \dots, M \quad (20)$$

where  $f_{\text{dec}}$  denotes the decoder, and  $M (M \geq 1)$  is the number of time steps to predict. This means that errors in the prediction process accumulate over  $M$  steps. As the decoder of NRI only uses the current state to predict future states, it is difficult to avoid error accumulation. Apparently, the current state of the interacting system is related to the previous states, and thus, incorporating historical information can help learn the dynamical rules and alleviate the error accumulation in multi-step prediction. To this end, sequence models are used to capture the non-linear correlation between previous states and the current state. GNNs combined with sequence models can be used to capture spatio-temporal information, which has been widely used in traffic flow forecasting (Li et al. 2018; Guo et al. 2019). In this paper, the sequence model can be composed of one or more of CNNs, RNNs, attention mechanisms, etc.

Inspired by interaction networks (Battaglia et al. 2016) and graph networks (Sanchez-Gonzalez et al. 2018), sequence models are added to the original node-to-edge and edge-to-node operations to obtain their spatio-temporal versions. In this way, the decoder can integrate the spatio-temporal interaction information at different fine-grained levels, i.e., the node level and the edge level, which can help better learn the dynamical rules.

As shown in Fig. 4, the decoder contains a node-to-edge and an edge-to-node spatio-temporal message passing operations. The node-to-edge spatio-temporal message passing operation is defined as

$$v \rightarrow e: \quad \hat{\mathbf{e}}_{(i,j)}^t = \sum_k z_{ij,k} \hat{f}_e^k([\mathbf{x}_i^t, \mathbf{x}_j^t]), \quad (21)$$

$$\hat{\mathbf{h}}_{(i,j)}^{t+1} = S_{\text{edge}}([\hat{\mathbf{e}}_{(i,j)}^t, \hat{\mathbf{h}}_{(i,j)}^{1:t}]), \quad (22)$$

where  $z_{ij,k}$  is the  $k$ -th element of the vector  $\mathbf{z}_{ij}$ , representing the edge type  $k$ , whose effect is modeled by an MLP  $\hat{f}_e^k$ . For each edge  $e_{ij}$ , the effect of all potential interactions are aggregated into  $\hat{\mathbf{e}}_{(i,j)}^t$  as a weighted sum. Its concatenation with the previous hidden states  $\hat{\mathbf{h}}_{(i,j)}^{1:t}$  is fed to  $S_{\text{edge}}$  to generate the future hidden state of interactions  $\hat{\mathbf{h}}_{(i,j)}^{t+1}$  that captures the temporal dependencies at edge level.

Similarly, the edge-to-node spatio-temporal message

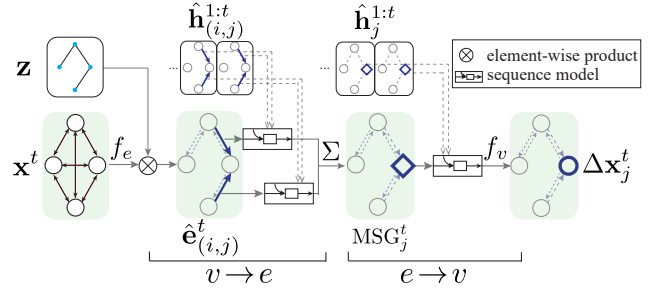


Figure 4: The structure of the decoder. The decoder takes the interacting relations  $\mathbf{z}$ , the current state  $\mathbf{x}^t$ , the historical hidden states  $\hat{\mathbf{h}}_{(i,j)}^{1:t}$  and  $\hat{\mathbf{h}}_j^{1:t}$  as inputs to predict the change in state  $\Delta \mathbf{x}^t$ . Sequence models are added to the message passing operations to jointly capture the spatio-temporal dependencies. Elements that are currently updated are highlighted in blue, e.g.,  $\hat{\mathbf{e}}_{(i,j)}^t$  is an edge embedding updated in the node-to-edge ( $v \rightarrow e$ ) message passing operation.

passing operation is defined as

$$e \rightarrow v: \text{MSG}_j^t = \sum_{i \neq j} \hat{\mathbf{h}}_{(i,j)}^{t+1}, \quad (23)$$

$$\hat{\mathbf{h}}_j^{t+1} = S_{\text{node}}([\text{MSG}_j^t, \mathbf{x}_j^t, \hat{\mathbf{h}}_j^{1:t}]). \quad (24)$$

For each node  $v_j$ , the spatial dependencies  $\hat{\mathbf{h}}_{(i,j)}^{t+1}$  are aggregated into  $\text{MSG}_j^t$ .  $\text{MSG}_j^t$  together with the current state  $\mathbf{x}_j^t$  and the previous hidden states  $\hat{\mathbf{h}}_j^{1:t}$  are fed to  $S_{\text{node}}$  to generate the future hidden state of nodes  $\hat{\mathbf{h}}_j^{t+1}$  that captures the temporal dependencies at node level.

Finally, the predicted future state is defined as

$$\Delta \mathbf{x}_j^t = \hat{f}_v(\hat{\mathbf{h}}_j^{t+1}), \quad (25)$$

$$\mu_j^{t+1} = \mathbf{x}_j^t + \Delta \mathbf{x}_j^t, \quad (26)$$

$$p_{\theta}(\mathbf{x}^{t+1} | \mathbf{x}^{1:t}, \mathbf{z}) = \mathcal{N}(\mathbf{x}^{t+1} | \mu^{t+1}, \sigma^2 \mathbf{I}), \quad (27)$$

where  $\hat{f}_v$  is an MLP predicting the change in state  $\Delta \mathbf{x}_j^t$ , and  $\sigma^2$  is a fixed variance. Note that since both  $\hat{\mathbf{h}}_{(i,j)}^{1:t}$  and  $\hat{\mathbf{h}}_j^{1:t}$  rely on the historical states  $\mathbf{x}^{1:t}$ , the decoder is implicitly a function of  $\mathbf{x}^{1:t}$ .

### Structural Prior Knowledge

As the scale and complexity of systems increase, it becomes more difficult to infer relations solely based on the state sequences of individual agents. Therefore, it is desirable to incorporate possible structural prior knowledge, such as sparsity (Kipf et al. 2018), symmetry and node degree distribution (Li et al. 2019). Since symmetric relations widely exist in physical dynamical systems, the symmetry of relations is studied as a special case. Li et al. (2019) impose a hard constraint on symmetry, i.e., setting  $\mathbf{z}_{ji} = \mathbf{z}_{ij}$ . However, the hard constraint may limit the exploration of the model during the training procedure, and sometimes it may lead to a decrease of the prediction precision. By contrast, this paper

imposes a soft constraint by adding a regularization term to the original loss function.

Specifically, an auxiliary distribution  $q'_\phi(\mathbf{z}|\mathbf{x})$  is introduced as the “transpose” of the predicted relation distribution  $q_\phi(\mathbf{z}|\mathbf{x})$ , namely,

$$q'_\phi(\mathbf{z}_{ij}|\mathbf{x}) = q_\phi(\mathbf{z}_{ji}|\mathbf{x}). \quad (28)$$

Then, the Kullback-Leibler divergence between  $q'_\phi(\mathbf{z}|\mathbf{x})$  and  $q_\phi(\mathbf{z}|\mathbf{x})$  is used as a regularization term for symmetry, i.e.,

$$\mathcal{L}' = \mathcal{L} - \lambda \cdot \text{KL} [q'_\phi(\mathbf{z}|\mathbf{x})||q_\phi(\mathbf{z}|\mathbf{x})], \quad (29)$$

where  $\lambda$  is a penalty factor for the symmetric prior.

Notations used in this paper, details of the computation of  $\mathcal{L}'$  and the pseudo code of NRI-MPM are shown in the supplementary material<sup>1</sup>.

## Experiments

All methods are tested on three types of simulated physical systems: particles connected by springs, charged particles and phase-coupled oscillators, named Springs, Charged and Kuramoto, respectively. For the Springs and Kuramoto datasets, objects do or do not interact with equal probability. For the Charged datasets, objects attract or repel with equal probability. For each type of system, a 5-object dataset and a 10-object dataset are simulated. All datasets contain 50k training samples, and 10k validation and test samples. Further details on data generation can be found in (Kipf et al. 2018). All methods are evaluated w.r.t. two metrics, the accuracy of relation reconstruction and the mean squared error (MSE) of future state prediction.

For our method, the sequence models used in both encoders and decoders are composed of gated recurrent units (GRUs) (Cho et al. 2014), except that attention mechanisms are added to the decoder in the Kuramoto datasets. The penalty factor  $\lambda$  is set to  $10^2$  for all datasets except that it is set to 1 and  $10^3$  for the 5-object and 10-object Kuramoto datasets, respectively. Details of experimental settings are shown in the supplementary material.

### Baselines

To evaluate the performance of our method, we compare it with several competitive methods as follows.

- Correlation (Kipf et al. 2018): a baseline that predicts the relations between two particles based on the correlation of their state sequences.
- LSTM (Kipf et al. 2018): an LSTM that takes the concatenation of all state vectors and predict all future states simultaneously.
- NRI (Kipf et al. 2018): the neural relational inference model that jointly learns the relations and the dynamics with a VAE.
- SUGAR (Li et al. 2019): a method that introduces structural prior knowledge such as hard symmetric constraint and node degree distribution for relational inference.

<sup>1</sup><https://arxiv.org/pdf/2101.09486>

Model	Springs	Charged	Kuramoto
5 objects			
Correlation	52.4 ± 0.0	55.8 ± 0.0	62.8 ± 0.0
NRI	<b>99.9 ± 0.0</b>	82.1 ± 0.6	96.0 ± 0.1
SUGAR	<b>99.9 ± 0.0*</b>	82.9 ± 0.8*	91.8 ± 0.1*
ModularMeta	<b>99.9</b>	88.4	96.2 ± 0.3*
NRI-MPM	<b>99.9 ± 0.0</b>	<b>93.3 ± 0.5</b>	<b>97.3 ± 0.2</b>
Supervised	99.9 ± 0.0	95.4 ± 0.1	99.3 ± 0.0
10 objects			
Correlation	50.4 ± 0.0	51.4 ± 0.0	59.3 ± 0.0
NRI	98.4 ± 0.0	70.8 ± 0.4	75.7 ± 0.3
SUGAR	98.3 ± 0.0*	72.0 ± 0.9*	74.0 ± 0.2*
ModularMeta	98.8 ± 0.0*	63.8 ± 0.1*	<b>89.6 ± 0.1*</b>
NRI-MPM	<b>99.1 ± 0.0</b>	<b>81.6 ± 0.2</b>	80.3 ± 0.6
Supervised	99.4 ± 0.0	89.7 ± 0.1	94.3 ± 0.8

\* Results in these datasets are unavailable in the original papers. Results of ModularMeta are obtained by running the codes provided by the authors, and those of SUGAR are reproduced by ourselves.

Table 1. Accuracy (%) of relation reconstruction.

- ModularMeta (Alet et al. 2019): a method that solves the relational inference problem via modular meta-learning.

To compare the performance of our method with “gold standard” methods, i.e., those trained given the ground truth relations, this paper introduces two variants as follows.

- Supervised: a variant of our method that only trains an encoder with the ground truth relations.
- NRI-MPM (true graph): a variant of our method that only trains a decoder given the ground truth relations.

The codes of NRI<sup>2</sup> and ModularMeta<sup>3</sup> are public and thus directly used in our experiments. SUGAR is coded by ourselves according to the original paper. Note that Correlation and Supervised are only designed for relation reconstruction, while LSTM and NRI-MPM (true graph) are only designed for state prediction.

To verify the effectiveness of the proposed message passing mechanisms and the structural prior knowledge, this paper introduces some variants of our method as follows.

- NRI-MPM w/o RI: a variant of our method without the relation interaction mechanism.
- NRI-MPM w/o intra-RI, NRI-MPM w/o inter-RI: variants of our method without the intra- and inter-edge interaction operations, respectively.
- NRI-MPM w/o ST: a variant of our method without the spatio-temporal message passing mechanism.
- NRI-MPM w/o Sym, NRI-MPM w/ hard Sym: variants of our method without the symmetric prior and that with hard symmetric constraint, respectively.

### Comparisons with Baselines

The results of relation reconstruction and future state prediction in the 5-object datasets (for MSEs in the 10-object

<sup>2</sup><https://github.com/ethanfetaya/nri>

<sup>3</sup><https://github.com/FerranAlet/modular-metalearning>

Datasets	Springs			Charged			Kuramoto		
Predictions steps	1	10	20	1	10	20	1	10	20
LSTM	4.13e-8	2.19e-5	7.02e-4	1.68e-3	6.45e-3	1.49e-2	<b>3.44e-4</b>	<b>1.29e-2</b>	4.74e-2
NRI	3.12e-8	3.29e-6	2.13e-5	1.05e-3	3.21e-3	7.06e-3	1.40e-2	2.01e-2	<b>3.26e-2</b>
SUGAR	3.71e-8*	3.86e-6*	1.53e-5*	1.18e-3*	3.43e-3*	7.38e-3*	2.12e-2*	9.45e-2*	1.83e-1*
ModularMeta	3.13e-8	3.25e-6	-	1.03e-3	3.11e-3	-	2.35e-2*	1.10e-1*	1.96e-1*
NRI-MPM	<b>8.89e-9</b>	<b>5.99e-7</b>	<b>2.52e-6</b>	<b>7.29e-4</b>	<b>2.57e-3</b>	<b>5.41e-3</b>	1.57e-2	2.73e-2	5.36e-2
NRI-MPM (true graph)	1.60e-9	9.06e-9	1.50e-7	8.06e-4	2.51e-3	5.66e-3	1.73e-2	2.49e-2	4.09e-2

\* Results in these datasets are unavailable in the original papers. Results of ModularMeta are obtained by running the codes provided by the authors, and those of SUGAR are reproduced by ourselves.

Table 2. Mean squared error in predicting future states for simulations with 5 interacting objects.

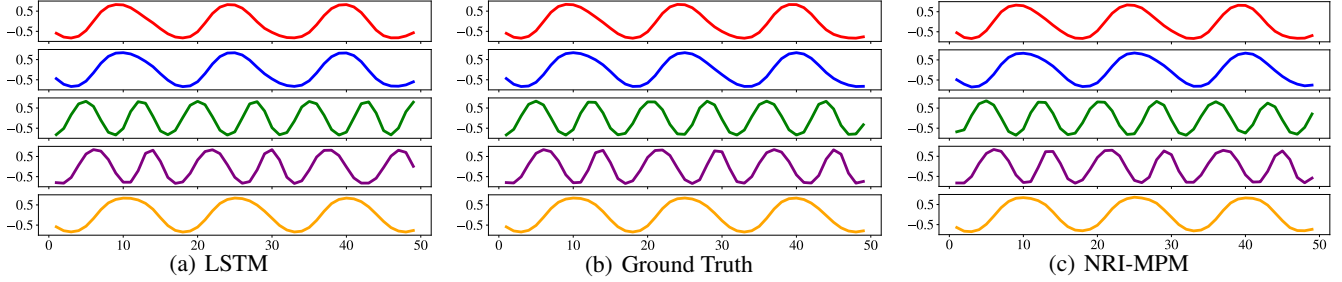


Figure 5: Visualization of the ground truth states (b) together with the predicted states for LSTM (a) and NRI-MPM (c) in the 5-object Kuramoto dataset.

datasets, see the supplementary material) are shown in Table 1 and Table 2, respectively. Our method significantly outperforms all baselines in nearly all datasets in terms of both accuracy and MSE. In the Springs datasets, the accuracies of NRI, SUGAR, ModularMeta and our method are all comparable with the supervised baseline, while our method achieves significantly lower MSEs in the 5-object systems. This indicates that our method can better learn the relations and dynamics simultaneously. In the Charged datasets, our method outperforms the baselines by 4.9%-9.6% in terms of accuracy. The reason may be that the charged systems are densely connected since a charged particle interacts with all other particles, and our method can better handle this situation. SUGAR achieves higher accuracies than NRI, suggesting the effectiveness of structural prior knowledge. Still, our method performs better with the extra help of the proposed message passing mechanisms.

In the Kuramoto datasets, the performance of our method is non-dominant. Our method achieves higher accuracies than NRI, while the MSEs are larger. ModularMeta achieves higher accuracy than our method in the 10-object system, while the MSEs are much poorer than all other methods. Maybe the interactions among objects in the Kuramoto dataset are relatively weak (Kuramoto 1975), making it more difficult to infer the relations based on the observed states, and the situation worsens as the scale of systems increases. ModularMeta infers all relations as a whole, and modifies its predictions with the help of simulated annealing, which may help better search the structural space to achieve better relation predictions. However, this does not translate to

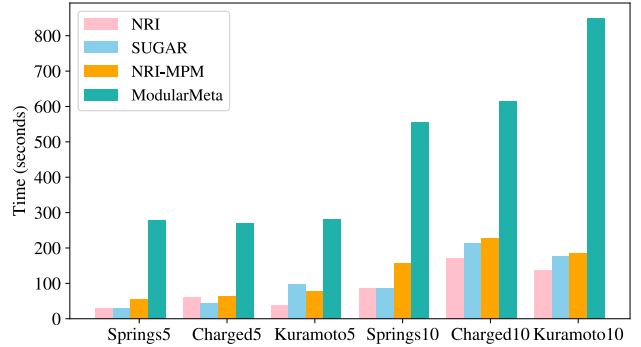


Figure 6: Running times of different methods.

better state prediction. According to (Kuramoto 1975), Kuramoto is a synchronization system where the object states converge to a certain set of values in the long run. Maybe the interaction is less helpful for state prediction in this case.

**Qualitative Analysis of Future State Prediction** One observes that LSTM achieves lower MSEs for short-term prediction in the 5-object Kuramoto dataset, but its performance declines for long-term prediction. To further understand the predictive behavior of LSTM and our method, we conduct a qualitative analysis by visualizing the predicted states together with the ground truth states in 49 steps, and the results are shown in Fig. 5. From Fig. 5(a), LSTM can capture the shape of the sinusoidal waveform but fails to make accurate

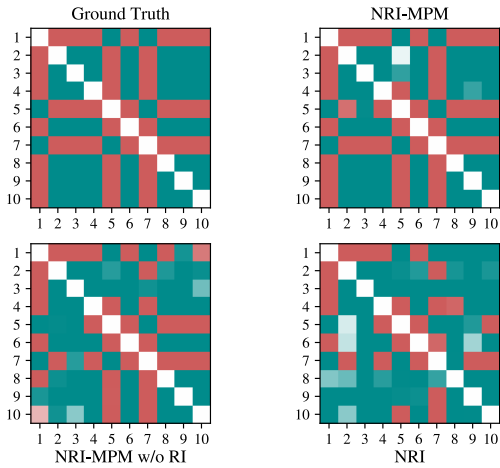


Figure 7: Predicted relations of different methods.

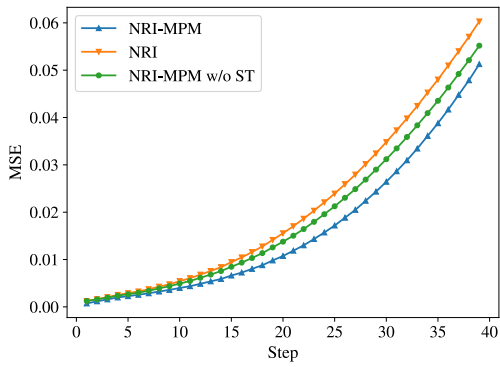


Figure 8: MSEs of multi-step prediction.

prediction for time steps larger than 40 (e.g., curves in green and purple). By contrast, as shown in Fig. 5(c), the predicted states of our method closely match the ground truth states except for those in the last few time steps for the third particle, whose curve is colored in green. Maybe our method can better capture the interactions among particles that affect the long-term dynamics. Note that this result is consistent with that reported in *Appendix A.1* in the original paper of NRI.

**Running Time of Different Methods** The running times of different methods in a single epoch are reported in Fig. 6. It can be seen that our method requires some more time than NRI, which is a natural consequence of more complex models. The running time of SUGAR is comparable with our method. It is worth noting that ModularMeta requires much more running time than the others and the situation becomes more severe in larger systems. Maybe meta-learning the proposal function is computationally costly in ModularMeta. These results show that our method can achieve better performance with a smaller additional cost of computation.

### Ablation Study

The ablation studies are conducted in the 10-object Charged dataset, and the results are shown in Table 3.

Metrics	Accuracy	MSEs		
Prediction steps		1	10	20
NRI-MPM w/o RI	$77.5 \pm 0.2$	$8.07e-4$	$4.13e-3$	$1.13e-2$
NRI-MPM w/o intra-RI	$80.5 \pm 0.4$	$7.83e-4$	$3.90e-3$	$1.26e-2$
NRI-MPM w/o inter-RI	$78.3 \pm 0.7$	$7.71e-4$	$4.11e-3$	$1.13e-2$
NRI-MPM w/o ST	$74.0 \pm 0.5$	$1.26e-3$	$4.92e-3$	$1.38e-2$
NRI-MPM w/o Sym	$79.3 \pm 0.7$	$7.60e-4$	<b><math>3.77e-3</math></b>	$1.06e-2$
NRI-MPM w/ hard Sym	$80.4 \pm 0.5$	$7.96e-4$	$3.83e-3$	$1.04e-2$
NRI-MPM	<b><math>81.6 \pm 0.2</math></b>	<b><math>7.52e-4</math></b>	$3.80e-3$	<b><math>1.03e-2</math></b>

Table 3. Ablation study in the 10-object Charged dataset.

**Effect of Relation Interaction Mechanism** As shown in Table 3, the accuracy drops significantly by removing the relation interacting mechanism, verifying its effectiveness in relation prediction. Besides, the MSEs decrease, indicating that more accurate relation prediction can help with future state prediction. Besides, the contribution of the inter-edge interaction is higher than that of the intra-edge interaction. Intuitively, the intra- and inter-edge operations capture local and global interactions, respectively, and maybe in this dataset, global interactions are more informative than local interactions.

To gain an intuitive understanding of the effect of the relation interaction mechanism, we conduct a case study on the 10-object charged particle systems. The distributions of predicted relations of NRI-MPM, NRI-MPM w/o RI and NRI together with the ground truth relations are visualized in Fig. 7. The two types of relations are highlighted in red and green, respectively, while the diagonal elements are all in white as there no self-loops. It is known that two particles with the same charge repel each other while those with the opposite charge attract each other. Consequently, for any particles  $v_i, v_j$ , and  $v_k$ , the relations  $e_{ij}$  and  $e_{ik}$  are correlated. As shown in Fig. 7, our method can model the dependencies among all relations much better than NRI, while removing the relation interaction mechanism results in less consistent prediction.

### Effect of Spatio-temporal Message Passing Mechanism

As shown in Table 3, the MSEs increase by removing the spatio-temporal message passing mechanism, and the differences are more significant for  $M = 10$  and  $M = 20$ , verifying that the spatio-temporal message passing mechanism can alleviate error accumulation in multi-step prediction. Furthermore, the MSEs of different methods that predict 40 steps into the future are shown in Fig. 8. Note that the differences among different methods narrow for large  $M$ . Maybe it is still challenging for all methods to handle error accumulation in the long run.

Interestingly, the decrease of NRI-MPM w/o ST in terms of accuracy is more significant than that of NRI-MPM w/o RI. Maybe the spatio-temporal message passing mechanism imposes strong dependencies of historical interactions, which indirectly helps with relation reconstruction.

**Effect of Symmetric Prior** As shown in Table 3, without the symmetric prior, the accuracy of NRI-MPM decreases by

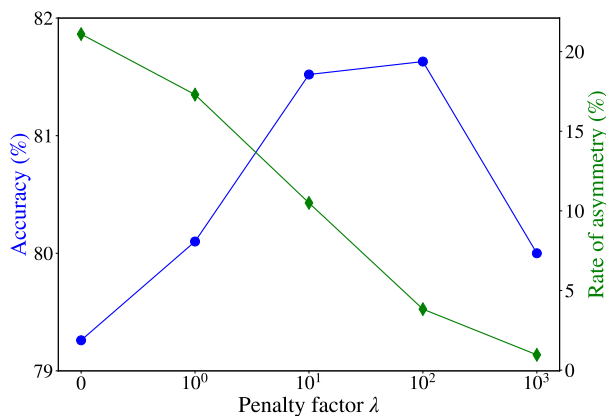


Figure 9: The accuracy and the rate of asymmetry w.r.t.  $\lambda$ .

2.3%, while the MSEs are on par with the original model, indicating that the symmetric prior can help with relation reconstruction without hurting the precision of future state prediction. Compared with NRI-MPM w/ hard Sym, our method achieves higher accuracy with lower MSEs. Maybe the hard constraint of symmetry limits the exploration of the model in the training procedure, while a soft constraint provides more flexibility.

The effect of the penalty factor  $\lambda$  is shown in Fig. 9. The rate of asymmetry decreases significantly as  $\lambda$  increases, while the accuracy increases steadily and peaks around  $\lambda = 10^2$ , verifying that adjusting  $\lambda$  can control the effect of the symmetric prior and reasonable values will benefit relation reconstruction.

## Related Work

This paper is part of an emerging direction of research attempting to model the explicit interactions in dynamical systems using neural networks as in NRI (Kipf et al. 2018).

Most closely related are the papers of Alet et al. (2019), Li et al. (2019), Webb et al. (2019) and Zhang et al. (2019). Alet et al. (2019) frame this problem as a modular meta-learning problem to jointly infer the relations and use the data more effectively. To deal with more complex systems, Li et al. (2019) incorporate various structural prior knowledge as a complement to the observed states of agents. Webb et al. (2019) extend NRI to multiplex interaction graphs. Zhang et al. (2019) explore relational inference over a wider class of dynamical systems, such as discrete systems and chaotic systems, assuming a shared interaction graph for all state sequences, which is different from the experimental settings of Kipf et al. (2018) and Alet et al. (2019). Compared with this line of work, our method focuses on introducing efficient message passing mechanisms to enrich the representative power of NRI.

Many recent works seek to extend the message passing mechanisms of GNNs. Zhu et al. (2020) define a bilinear aggregator to incorporate the possible interactions among all neighbors of a given node. Brockschmidt (2020) defines node aware transformations over messages to impose feature-wise modulation. Nevertheless, these variants

treat the messages as a transient part of the computation of the node embeddings, our relation interaction mechanism is aimed at learning edge embeddings. Herzig et al. (2019) use non-local operations to capture the interactions among all relations, which requires quadratic time complexity.

Besides, many works extend GNNs to handle structured time series. Graph convolutional networks with CNNs (Yu, Yin, and Zhu 2018), GRUs (Li et al. 2018), or attention mechanisms (Guo et al. 2019) are introduced to deal with spatial and temporal dependencies separately for traffic flow forecasting. Furthermore, spatio-temporal graphs (Song et al. 2020) and spatio-temporal attention mechanisms (Zheng et al. 2020) are proposed to capture complex spatio-temporal correlations. Our methods borrow these ideas from traffic flow forecasting to define spatio-temporal message passing operations for neural relational inference.

## Conclusion

This paper introduces efficient message passing mechanisms with structural prior knowledge for neural relational inference. The relation interaction mechanism can effectively capture the coexistence of all relations and help make a joint prediction. By incorporating historical information, the spatio-temporal message passing mechanism can effectively alleviate error accumulation in multi-step state prediction. Additionally, the structural prior knowledge, symmetry as a special case, can promote the accuracy of relation reconstruction in more complex systems. The results of extensive experiments on simulated physics systems validate the effectiveness of our method.

Currently, only simple yet effective implementations using GRUs and attention mechanisms are adopted for the sequence models. Future work includes introducing more advanced models like the Transformers to further improve the performance. Besides, current experiments are conducted on simulated systems over static and homogeneous graphs. Future work includes extending our method to systems over dynamic and heterogeneous graphs. Furthermore, the proposed method will be applied to study the mechanism of the emergence of intelligence in different complex systems, including multi-agent systems, swarm systems, physical systems and social systems.

## Acknowledgements

This work is supported by the National Key R&D Program of China (2018AAA0101203), and the National Natural Science Foundation of China (62072483, 61673403, U1611262). This work is also supported by MindSpore.

## References

- Alet, F.; Weng, E.; Lozano-Pérez, T.; and Kaelbling, L. P. 2019. Neural Relational Inference with Fast Modular Meta-learning. In *NeurIPS*, 11827–11838.
- Almaatouq, A.; Noriega-Campero, A.; Alotaibi, A.; Krafft, P. M.; Moussaid, M.; and Pentland, A. 2020. Adaptive Social Networks Promote the Wisdom of Crowds. *PNAS* 117(21): 11379–11386.



- Bapst, V.; Keck, T.; Grabska-Barwińska, A.; Donner, C.; Cubuk, E. D.; Schoenholz, S. S.; Obika, A.; Nelson, A. W. R.; Back, T.; Hassabis, D.; and Kohli, P. 2020. Unveiling the Predictive Power of Static Structure in Glassy Systems. *Nature Physics* 16(4): 448–454.
- Battaglia, P.; Pascanu, R.; Lai, M.; Rezende, D. J.; and kavukcuoglu, K. 2016. Interaction Networks for Learning about Objects, Relations and Physics. In *NeurIPS*, 4509–4517.
- Brockschmidt, M. 2020. GNN-FiLM: Graph Neural Networks with Feature-wise Linear Modulation. In *ICML*.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations Using RNN Encoder–Decoder for Statistical Machine Translation. In *EMNLP*, 1724–1734.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural Message Passing for Quantum Chemistry. In *ICML*, 1263–1272.
- Guo, S.; Lin, Y.; Feng, N.; Song, C.; and Wan, H. 2019. Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting. In *AAAI*, 922–929.
- Ha, S.; and Jeong, H. 2020. Towards Automated Statistical Physics: Data-Driven Modeling of Complex Systems with Deep Learning. *arXiv preprint arXiv:2001.02539*.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive Representation Learning on Large Graphs. In *NeurIPS*, 1024–1034.
- Herzig, R.; Levi, E.; Xu, H.; Gao, H.; Brosh, E.; Wang, X.; Globerson, A.; and Darrell, T. 2019. Spatio-Temporal Action Graph Networks. In *ICCVW*.
- Hoshen, Y. 2017. VAIN: Attentional Multi-agent Predictive Modeling. In *NeurIPS*, 2701–2711.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- Kipf, T.; Fetaya, E.; Wang, K.-C.; Welling, M.; and Zemel, R. 2018. Neural Relational Inference for Interacting Systems. In *ICML*, 2688–2697.
- Kuramoto, Y. 1975. Self-Entrainment of A Population of Coupled Non-Linear Oscillators. In *International Symposium on Mathematical Problems in Theoretical Physics*, 420–422. Springer.
- Lee, J.; Lee, Y.; Kim, J.; Kosiorek, A.; Choi, S.; and Teh, Y. W. 2019. Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In *ICML*, 3744–3753.
- Li, J.; Yang, F.; Tomizuka, M.; and Choi, C. 2020. EvolveGraph: Multi-Agent Trajectory Prediction with Dynamic Relational Reasoning. In *NeurIPS*.
- Li, Y.; Meng, C.; Shahabi, C.; and Liu, Y. 2019. Structure-Informed Graph Auto-Encoder for Relational Inference and Simulation. In *ICML Workshop on Learning and Reasoning with Graph-Structured Representations*.
- Li, Y.; Yu, R.; Shahabi, C.; and Liu, Y. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *ICLR*.
- Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *ICLR*.
- Oliveira, M.; Pinheiro, D.; Macedo, M.; Bastos-Filho, C.; and Menezes, R. 2020. Uncovering the Social Interaction Network in Swarm Intelligence Algorithms. *Applied Network Science* 5(1): 1–20.
- Sanchez-Gonzalez, A.; Heess, N.; Springenberg, J. T.; Merel, J.; Riedmiller, M.; Hadsell, R.; and Battaglia, P. 2018. Graph Networks as Learnable Physics Engines for Inference and Control. In *ICML*, 4467–4476.
- Song, C.; Lin, Y.; Guo, S.; and Wan, H. 2020. Spatial-Temporal Synchronous Graph Convolutional Networks: A New Framework for Spatial-Temporal Network Data Forecasting. *AAAI* 914–921.
- van Steenkiste, S.; Chang, M.; Greff, K.; and Schmidhuber, J. 2018. Relational Neural Expectation Maximization: Unsupervised Discovery of Objects and their Interactions. In *ICLR*.
- Watters, N.; Zoran, D.; Weber, T.; Battaglia, P.; Pascanu, R.; and Tacchetti, A. 2017. Visual Interaction Networks: Learning A Physics Simulator from Video. In *NeurIPS*, 4539–4547.
- Webb, E.; Day, B.; Andres-Terre, H.; and Lió, P. 2019. Factorised Neural Relational Inference for Multi-Interaction Systems. *ICML Workshop on Learning and Reasoning with Graph-Structured Representations*.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks. In *ICLR*.
- Yang, Y.; Yu, L.; Bai, Y.; Wen, Y.; Zhang, W.; and Wang, J. 2018. A Study of AI Population Dynamics with Million-agent Reinforcement Learning. In *AAMAS*, 2133–2135.
- Yu, B.; Yin, H.; and Zhu, Z. 2018. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *IJCAI*, 3634–3640.
- Zhang, J.; Wang, W.; Xia, F.; Lin, Y.-R.; and Tong, H. 2020. Data-Driven Computational Social Science: A Survey. *Big Data Research* 100145.
- Zhang, Z.; Zhao, Y.; Liu, J.; Wang, S.; Tao, R.; Xin, R.; and Zhang, J. 2019. A General Deep Learning Framework for Network Reconstruction and Dynamics Learning. *Applied Network Science* 4(1): 1–17.
- Zheng, C.; Fan, X.; Wang, C.; and Qi, J. 2020. GMAN: A Graph Multi-Attention Network for Traffic Prediction. In *AAAI*, 1234–1241.
- Zhu, H.; Feng, F.; He, X.; Wang, X.; Li, Y.; Zheng, K.; and Zhang, Y. 2020. Bilinear Graph Neural Network with Neighbor Interactions. In *IJCAI*, 1452–1458.