

# Improving Temporal Knowledge Graph Forecasting via Multi-Rewards Mechanism and Confidence-Guided Tensor Decomposition Reinforcement Learning

Nam Le<sup>1,2,3</sup>, Thanh Le<sup>1,2,3</sup>, and Bac Le<sup>1,2,3</sup>

<sup>1</sup>Department of Computer Science

<sup>2</sup>Faculty of Information Technology, University of Science, Ho Chi Minh City, Vietnam

<sup>3</sup>Vietnam National University, Ho Chi Minh City, Vietnam

February 24, 2025

# Outline

Introduction

Backgrounds

Proposed Model

- Multi-reward Mechanism with Rule Enhancing

- Tensor Decomposition Confidence-Guided Policy Network

- Agent Parameter Learning

Experiments

Conclusion and Future Directions

# Introduction

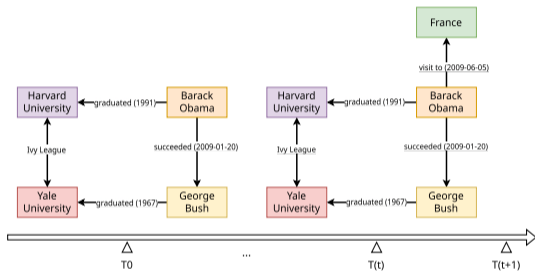


Figure: Example of Part of Temporal Knowledge Graph.

- ▶ In practice, data change over time.
- ▶ Reasoning problem on Temporal Knowledge Graph (TKG) can be viewed in two settings:
  - **Interpolation** which focuses on completing the missing links at past timestamps.
  - **Extrapolation** which focuses on forecasting future facts.

⇒ We mainly focus on **extrapolation** setting.

# Static & Temporal Knowledge Graph Reasoning

Table: Path-based Reasoning models for Static & Temporal Knowledge Graph

Method	Future timestamps	Unseen entities	Efficient	Explanatory	Reward Flexible	Action Selection	Open Source
MINERVA [20]	x	x	✓	x	x	x	✓
Multi-hop KG [19]	x	✓	✓	x	x	x	✓
RE-NET [16]	✓	x	x	x	x	x	✓
CyGNett [9]	✓	x	✓	x	x	x	✓
TANGO [6]	✓	x	x	x	x	x	✓
TAgent [8]	✓	x	x	x	x	x	x
TITer [7]	✓	✓	✓	✓	x	x	✓
TPath [3]	✓	x	✓	x	x	x	x
DREAM [2]	✓	x	x	x	✓	x	x
RLAT [1]	✓	x	x	x	x	x	x

## Challenges in RL for TKG Reasoning

General, when compared to interpolation, **extrapolation** setting is more difficult and challenging. Recently, path-based reasoning methods are potential solutions for this setting and face two challenges:

- ▶ The reward function is a critical component for the agent. **Most current works focus on constructing a binary global reward function, which makes the agent's learning process inflexible.**
- ▶ **The action space for the agent is too large, and there is limited research on how to select appropriate actions for the agent.**

## Our contributions

- ▶ Proposing a new multi-reward function, incorporating various reward criteria for the agent.
- ▶ Incorporating Tensor decomposition architectures such as TuckER, ComplEx, and LowFER with MLP and KAN-Policy Network to generate reliability scores for actions.
- ▶ Performing experiments and ablation study on standard datasets for the future link prediction task with improvements on metrics.

# Outline

Introduction

Backgrounds

Proposed Model

- Multi-reward Mechanism with Rule Enhancing

- Tensor Decomposition Confidence-Guided Policy Network

- Agent Parameter Learning

Experiments

Conclusion and Future Directions

## Basic Notations

- ▶ Let  $\mathcal{E}$ ,  $\mathcal{R}$ ,  $\mathcal{T}$ , and  $\mathcal{Q}$  denote the sets of entities, relations, timestamps and quadruples.
- ▶ Each quadruplet in TKG can be defined as a tuple  $(e_s, r, e_o, t)$ .



## Problem Statement

Considering TKG as  $\mathcal{G}_{(1,T)} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_T\}$ , where  $\mathcal{G}_t = \{\mathcal{E}_t, \mathcal{R}, \mathcal{Q}_t\}$  is a static multi-relational graph, and  $\mathcal{E}_t$  and  $\mathcal{Q}_t$  denote entities and facts that exist at time  $t$ .

- Input: with given a query  $(e_q, r_q, ?, t_q)$  or  $(?, r_q, e_q, t_q)$ , and a set of known facts  $\{(e_{s_i}, r_i, e_{o_i}, t_i) | t_i < t_q\}$
- Output: potential candidates which can replace the missing object or subject entity in the input query.

## RL Framework for TKG reasoning

Mains components in framework are:

- States: Let  $\mathcal{S}$  as state space, each state can represent as  $s_\ell = (e_{(\ell)}, t_{(\ell)}, e_q, t_q, r_q) \in \mathcal{S}$ .
- Actions: Let  $\mathcal{A}$  be the action space. Set of actions for step  $\ell$  is  $\mathcal{A}_\ell = \{(r', e', t') | (e_\ell, r', e', t') \in \mathcal{Q}, t' \leq t_\ell, t' < t_q\}$  which implies outgoing edges of the current node of agent.
- Transition function  $\xi : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  defined by:

$$(s_\ell, \mathcal{A}_\ell) \mapsto (e_{\ell+1}, t_{\ell+1}, e_q, t_q, r_q) = s_{\ell+1} \quad (1)$$

which transfer the environment state to a new node through edge selected by agent.

- Reward function: Commonly, binary global reward function is defined by:

$$R_{bin}(s_L) = \mathbb{I}(e_\ell == e_{gt}), \quad (2)$$

where  $\mathbb{I}(\cdot)$  is a function that return 1 or 0.

# Outline

Introduction

Backgrounds

**Proposed Model**

Multi-reward Mechanism with Rule Enhancing

Tensor Decomposition Confidence-Guided Policy Network

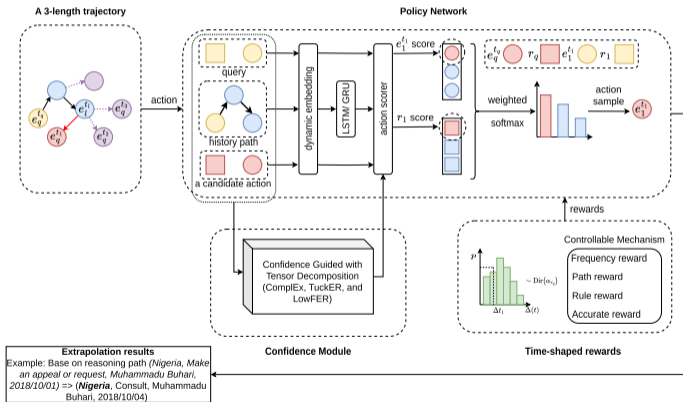
Agent Parameter Learning

Experiments

Conclusion and Future Directions

# Overview of our proposed model CATTer

Inspired by path-based methods for static & temporal KGs, we propose new temporal path-based reinforcement learning for extrapolated TKG reasoning with two advances: 1) multi-reward function and 2) confidence-guided policy network.



## Overview of our proposed model CATTer

Following TITer, there is no edge between snapshots, so we add three types of edges:

1. Reversed edges
2. Self-loop edges
3. Temporal Edges

## Multi-reward Mechanism with Rule Enhancing

We proposed new multi-reward function:

$$R = (1 + \alpha_1 R_{\text{gt}})(1 + \alpha_2 R_{\text{rule}})(R_{\text{bin}} + \alpha_3 R_{\text{path}}), \quad (3)$$

where  $R_{\text{bin}}$  is binary reward,  $R_{\text{gt}}$  is (adjusted) ground truth frequency reward,  $R_{\text{rule}}$  is high-frequency rule reward,  $R_{\text{path}}$  is (adjusted) path length reward, and  $\alpha_1 \in (0, 1)$ ,  $\alpha_2 \in (0, 1)$ , and  $\alpha_3 \in (0, 1)$  are weights.

## Binary global reward

The *binary global reward* that is defined by:

$$R_{bin}(s_L) = \mathbb{I}(e_\ell == e_{gt}). \quad (4)$$

## Adjusted ground truth frequency reward

With given  $(e_q, r_q, e_{gt}, t_q)$ ,  $N_{gt} = \{n_1, n_2, \dots, n_m\}$  denote the number of times that the  $e_{gt}$  occur in  $m$  snapshot  $\{G_{t_q-1}, G_{t_q-2}, \dots, G_{t_q-m}\}$ , i.e.,  $n_i$ , ( $i = 1, \dots, m$ ) is the number of times that  $e_{gt}$  occurs in subgraph  $G_{t_q-i}$ . We expect the  $e_{gt}$  should occur maximum as possible

We define the ground truth frequency reward as follows:

$$R_{gt}(s_L) = \begin{cases} f_i, & \text{if } t_{q-m} \leq t_i \leq t_q, \\ 0, & \end{cases} \quad (5)$$

where

$$f_i = \frac{n_i}{\max(N_{gt}) - \min(N_{gt})}. \quad (\text{normalized by max and min})$$



## Adjusted path length reward

We expect the path length to  $e_{gt}$  should be minimum as possible. So, we proposed *adjusted path length reward* which can be defined as:

$$R_{\text{path}}(s_L) = \frac{w_{\text{path}}}{p_\ell - 1} \quad (6)$$

where  $p_\ell \leq p_{\text{max}}$  denotes the length of the path taken by the agent to capture the target entity from the source node at step  $\ell$ ,  $p_{\text{max}}$  is the maximum path length which agent can reach a node, and  $w_{\text{path}} \in (0, 1)$  is the weight for current path length which is taken.

Note that: minus one in denominator means to accelerate our expectations.

## High-frequency rule reward

In our observations, knowledge graphs usually contain a pair entity relation, frequently appearing in the timelines.

Formally, given a common pair entity-relation set, which is denoted as  $ER = \{(e_i, r_i)\}_{i=1}^k$ . Each pair in ER has a frequency of occurrence greater than or equal to a threshold  $\vartheta$  depending on the dataset. Then, we define a *high-frequency rule reward* for our agent as follows:

$$R_{\text{rule}}(s_L) = \begin{cases} w_{\text{rule}}, & \text{if } (e_\ell, r_\ell) \in ER, \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where  $w_{\text{rule}}$  is reward value for matching rule.

## Multi-reward reshaping

Based on the training set, we estimate a [Dirichlet Distribution](#) for each relation. Then, we reshape the original multi-reward with this distribution:

$$\begin{aligned}\tilde{R}(s_L) &= (1 + p_{\Delta t_L} R(s_L)) \\ \Delta t_L &= t_q - t_L \\ (p_1, \dots, p_k) &\sim \text{Dirichlet}(\boldsymbol{\alpha}_{r_q}), \boldsymbol{\alpha}_{r_q} \in \mathbb{R}^K\end{aligned}\tag{8}$$

## Policy Network

- ▶ **Dynamic Embedding.** Relative time encoding function to represent time information.

$$\Phi(t_q - t) = \sigma(\mathbf{w}\Delta t + \mathbf{b}) = \Phi(\Delta t); \quad (9)$$

$$\mathbf{e}_i^t = [\mathbf{e}_i; \Phi(\Delta t)] \quad (10)$$

- ▶ **Historical Path Encoding.** We use LSTM/ GRU to encode the search history which is the sequence of actions taken.

$$\begin{aligned} \mathbf{h}_\ell^{\text{gru}} &= \text{GRU}([\mathbf{r}_{\ell-1}; \mathbf{e}_{\ell-1}^{t_{\ell-1}}], \mathbf{h}_{\ell-1}), \\ \mathbf{h}_0^{\text{gru}} &= \text{GRU}([\mathbf{r}_0; \mathbf{e}_q^{t_q}, \mathbf{0}]). \end{aligned} \quad (11)$$

with  $\mathbf{r}_0$  is dummy relation for initialization. Similar if we use LSTM.

## Policy Network

- **Action scoring.** We use a weighted action scoring to help agent pay more attention to attributes of destination nodes.

$$\phi(a_n, s_\ell) = \beta_n \langle \tilde{\mathbf{e}}, \mathbf{e}_n^{t_n} \rangle + (1 - \beta_n) \langle \tilde{\mathbf{r}}, \mathbf{r}_n \rangle, \quad (12)$$

with

$$\begin{aligned} \tilde{\mathbf{e}} &= \mathbf{W}_e \text{ReLU}(\mathbf{W}_1[\mathbf{h}_\ell^{\text{lstm/gru}}; \mathbf{e}_q^{t_q}; \mathbf{r}_q]), \\ \tilde{\mathbf{r}} &= \mathbf{W}_r \text{ReLU}(\mathbf{W}_1[\mathbf{h}_\ell^{\text{lstm/gru}}; \mathbf{e}_q^{t_q}; \mathbf{r}_q]), \\ \beta_n &= \text{sigmoid}(\mathbf{W}_\beta[\mathbf{h}_\ell^{\text{lstm/gru}}; \mathbf{e}_q^{t_q}; \mathbf{r}_q; \mathbf{e}_n^{t_n}; \mathbf{r}_n]), \end{aligned}$$

where  $\mathbf{W}_1$ ,  $\mathbf{W}_e$ ,  $\mathbf{W}_r$  and  $\mathbf{W}_\beta$  are trainable parameters for MLP or KAN.

## Policy Network

- **Confidence Rate Action Calculation.** We calculate the confidence rate  $c_{a_n|q}$  of each  $a_n \in \mathcal{A}_\ell$  via softmax function which receive the input vector from tensor decomposition such as Tucker [15], Complex [22], and LowFER [11] as follow:

$$c_{a_n|q} = \frac{\exp(\psi_{a_n|q})}{\sum_{a'_\ell \in \mathcal{A}_\ell} \exp(\psi_{a'_\ell|q})}, \quad (13)$$

where

$$\psi_{a_n|q} = \mathcal{W} \times_1 \mathbf{e}_q^{t_q} \times_2 \mathbf{r}_q \times_3 \mathbf{e}_n^{t_n}, \text{ if use Tucker,}$$

$$\psi_{a_n|q} = \text{Re} \left( \left\langle \mathbf{e}_q^{t_q}, \mathbf{r}_q, \overline{\mathbf{e}_n^{t_n}} \right\rangle \right) \text{ if use Complex,}$$

$$\psi_{a_n|q} = (\mathbf{S}^k \text{diag}(\mathbf{U}^\top \mathbf{e}_q^{t_q}) \mathbf{V}^\top \mathbf{r}_q)^\top \mathbf{e}_n^{t_n}, \text{ if use LowFER,}$$

## Policy Network

- ▶ The policy  $\pi_\theta(a_\ell | s_\ell)$  at step  $\ell$  is defined as:

$$\pi_\theta(a_\ell | s_\ell) = \frac{\exp(\phi(a_\ell, s_\ell) * c_{a_\ell|q})}{\sum_{a'_\ell \in \mathcal{A}_\ell} \exp(\phi(a'_\ell, s_\ell) * c_{a'_\ell|q})} \quad (14)$$

## Optimization

We apply REINFORCE algorithm [27] that will iterate through all quadruple in  $\mathcal{Q}_{train}$  and update  $\theta$  with the following stochastic gradient method such as SGD [21], Adam [21, 24] or AdaGrad [26]:

$$\nabla_{\theta} J(\theta) \approx \nabla_{\theta} \sum_{m \in [1, L]} \tilde{R}(s_L | e_s, r, t) \log \pi_{\theta}(a_{\ell} | s_{\ell}) \quad (15)$$



# Outline

Introduction

Backgrounds

Proposed Model

- Multi-reward Mechanism with Rule Enhancing

- Tensor Decomposition Confidence-Guided Policy Network

- Agent Parameter Learning

Experiments

Conclusion and Future Directions

# Datasets and Baselines

- ▶ **Baselines.**

1. Interpolation-based models: TTransE [18], TA-DistMult [17], DE-Simple [12], and TNTCompEx [14].
2. Extrapolation-based models: RE-NET [13], CyGNet [10], TANGO [4], xERTE [5], and TITer [7].

- ▶ **Datasets.** ICEWS14 and ICEWS18 [23], WIKI [18] and YAGO [25].

## Performance and efficiency comparison

Method	ICEWS14				ICEWS18			
	MRR $\uparrow$	Hit@1 $\uparrow$	Hit@3 $\uparrow$	Hit@10	MRR $\uparrow$	Hit@1 $\uparrow$	Hit@3 $\uparrow$	Hit@10 $\uparrow$
TTransE	13.43	3.11	17.32	34.55	8.31	1.92	8.56	21.89
TA-DistMult	26.47	17.09	30.22	45.41	16.75	8.61	18.41	33.59
DE-SimplE	32.67	24.43	35.69	49.11	19.30	11.53	21.86	34.80
TNTComplEx	32.12	23.35	36.03	49.13	27.54	19.52	30.80	42.86
CyGNet	32.73	23.69	36.31	50.67	24.93	15.90	28.28	42.61
RE-NET	38.28	28.68	41.34	54.52	28.81	19.05	32.44	<b>47.51</b>
xERTE	40.79	32.70	45.67	57.30	29.31	21.03	<b>33.51</b>	46.48
TANGO-Tucker	–	–	–	–	28.68	19.35	32.17	47.04
TANGO-DistMult	–	–	–	–	26.75	17.92	30.08	44.09
TITer	<b>41.73</b>	<b>32.74</b>	<b>46.46</b>	<b>58.44</b>	<b>29.98</b>	<b>22.05</b>	33.46	44.83
TITer*	40.33	31.00	45.30	57.71	29.42	21.63	32.83	43.96
CATTer-MLP	<u>41.21</u>	<u>32.47</u>	<u>45.75</u>	<u>57.37</u>	<u>29.54</u>	<u>21.60</u>	32.99	44.51
CATTer-KAN	40.13	31.04	44.80	57.19	29.11	21.37	32.46	43.60
APG (%) $\uparrow$ (MLP)	<b>0.62</b>	<b>0.54</b>	<b>0.61</b>	<b>0.64</b>	<b>-0.22</b>	<b>-0.39</b>	<b>-0.02</b>	<b>0.03</b>
RPG (%) $\uparrow$ (MLP)	<b>2.18</b>	<b>4.74</b>	<b>0.99</b>	<b>-0.59</b>	<b>0.41</b>	<b>-0.14</b>	<b>0.49</b>	<b>1.25</b>
APG (%) $\uparrow$ (KAN)	<b>0.65</b>	<b>0.48</b>	<b>0.77</b>	<b>0.92</b>	<b>-0.61</b>	<b>-0.68</b>	<b>-0.54</b>	<b>-0.53</b>
RPG (%) $\uparrow$ (KAN)	<b>-0.49</b>	<b>0.13</b>	<b>-1.10</b>	<b>-0.90</b>	<b>-1.05</b>	<b>-1.20</b>	<b>-1.13</b>	<b>-0.82</b>

## Performance and efficiency comparison

Method	WIKI				YAGO			
	MRR $\uparrow$	Hit@1 $\uparrow$	Hit@3 $\uparrow$	Hit@10	MRR $\uparrow$	Hit@1 $\uparrow$	Hit@3 $\uparrow$	Hit@10 $\uparrow$
TTransE	29.27	21.67	34.43	42.39	31.19	18.12	40.91	51.21
TA-DistMult	44.53	39.92	48.73	51.71	54.92	48.15	59.61	66.71
DE-SimplE	45.43	42.6	47.71	49.55	54.91	51.64	57.30	60.17
TNTComplEx	45.03	40.04	49.31	52.03	57.98	52.92	61.33	66.69
CyGNet	33.89	29.06	36.10	41.86	52.07	45.36	56.12	63.77
RE-NET	49.66	46.88	51.19	53.48	58.02	53.06	61.08	66.29
xERTE	71.14	68.05	76.11	79.01	84.19	80.09	88.02	89.78
TANGO-Tucker	50.43	48.52	51.47	53.58	57.83	53.05	60.78	65.85
TANGO-DistMult	51.15	49.66	52.16	53.35	62.70	59.18	60.31	67.90
TITer	<b>75.50</b>	<b>72.96</b>	<b>77.49</b>	<b>79.02</b>	<u>87.47</u>	<u>84.89</u>	<b>89.96</b>	<u>90.27</u>
TITer*	73.56	71.48	74.86	76.40	87.80	85.52	89.92	90.31
CATTer-MLP	<u>74.18</u>	<u>72.02</u>	75.47	77.04	<b>87.58</b>	<b>85.13</b>	<u>89.90</u>	<b>90.34</b>
CATTer-KAN	74.21	71.96	75.63	77.32	87.19	84.84	89.38	89.78
APG (%) $\uparrow$ (MLP)	<b>0.88</b>	<b>1.47</b>	<b>0.45</b>	<b>-0.34</b>	<b>0.12</b>	<b>-0.03</b>	<b>0.16</b>	<b>0.55</b>
RPG (%) $\uparrow$ (MLP)	<b>0.84</b>	<b>0.76</b>	<b>0.81</b>	<b>0.83</b>	<b>-0.25</b>	<b>-0.46</b>	<b>-0.02</b>	<b>0.03</b>
APG (%) $\uparrow$ (KAN)	<b>-0.2</b>	<b>0.04</b>	<b>-0.5</b>	<b>-0.52</b>	<b>-0.31</b>	<b>-0.26</b>	<b>-0.37</b>	<b>-0.36</b>
RPG (%) $\uparrow$ (KAN)	<b>0.88</b>	<b>0.67</b>	<b>1.03</b>	<b>1.20</b>	<b>-0.69</b>	<b>-0.80</b>	<b>-0.60</b>	<b>-0.59</b>

## Performance and efficiency comparison

**Table:** Number of trainable parameters and calculation of our proposed models and baselines. MACs stand for Multi-Adds operations, and M stand for million.

Method	# Params	# MACs
RE-NET	5.459M	4.370M
CyGNet	8.568M	8.554M
xERTE	2.927M	225.895M
TITer	1.455M	0.225M
CATTer	1.425M	0.220M

# Convergence Study




Figure: The change of the loss function over each epoch with MLP and KAN Policy Network.

# Convergence Study

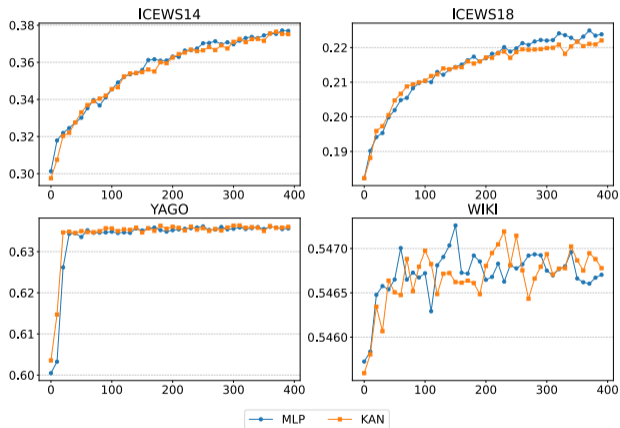



Figure: The change of the multi-reward function over each epoch with MLP and KAN Policy Network.  
Experiments

# The Effect of Tensor decomposition methods for action confidence





## The effect of multi-rewards




Figure: The effect of multi-reward mechanism for agent learning on ICEWS14, ICEWS18, YAGO and WIKI dataset.

# Outline

Introduction

Backgrounds

Proposed Model

- Multi-reward Mechanism with Rule Enhancing

- Tensor Decomposition Confidence-Guided Policy Network

- Agent Parameter Learning

Experiments

Conclusion and Future Directions

## Conclusion

- ▶ Introduces CATTer, an improved temporal-path-based RL model based on TimeTraveler.
- ▶ Integrates confidence probability into MLP and KAN layers.
- ▶ Designs a flexible Policy Network for better action selection.
- ▶ Employs a multi-reward function for improved adaptability in TKGs.
- ▶ Experimental results show enhanced future link prediction.
- ▶ Future work: Incorporating sub-graph patterns and temporal rules.

Thanks for your attention

## References I

- [1] Luyi Bai, Die Chai, and Lin Zhu. “RLAT: Multi-hop temporal knowledge graph reasoning based on Reinforcement Learning and Attention Mechanism”. In: *Knowledge-Based Systems* 269 (2023), p. 110514.
- [2] Shangfei Zheng et al. “DREAM: Adaptive Reinforcement Learning based on Attention Mechanism for Temporal Knowledge Graph Reasoning”. In: *arXiv preprint arXiv:2304.03984* (2023).
- [3] Luyi Bai et al. “Multi-hop reasoning over paths in temporal knowledge graphs using reinforcement learning”. In: *Applied Soft Computing* 103 (2021), p. 107144.
- [4] Zifeng Ding et al. “Temporal Knowledge Graph Forecasting with Neural ODE”. In: *arXiv preprint arXiv:2101.05151* (2021).
- [5] Zhen Han et al. “Explainable Subgraph Reasoning for Forecasting on Temporal Knowledge Graphs”. In: *International Conference on Learning Representations*. 2021.

## References II

- [6] Zhen Han et al. “Temporal knowledge graph forecasting with neural ode”. In: *arXiv preprint arXiv:2101.05151* (2021).
- [7] Haohai Sun et al. “TimeTraveler: Reinforcement Learning for Temporal Knowledge Graph Forecasting”. In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 2021, pp. 8306–8319.
- [8] Ye Tao, Ying Li, and Zhonghai Wu. “Temporal link prediction via reinforcement learning”. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2021, pp. 3470–3474.
- [9] Cunchao Zhu et al. “Learning from history: Modeling temporal knowledge graphs with sequential copy-generation networks”. In: *Proceedings of the AAAI conference on artificial intelligence*. 2021, pp. 4732–4740.
- [10] Cunchao Zhu et al. “Learning from History: Modeling Temporal Knowledge Graphs with Sequential Copy-Generation Networks”. In: *Thirty-Fifth AAAI Conference on Artificial Intelligence*. 2021, pp. 4732–4740.

## References III

- [11] Saadullah Amin et al. “LowFER: Low-rank bilinear pooling for link prediction”. In: *International Conference on Machine Learning*. PMLR. 2020, pp. 257–268.
- [12] Rishab Goel et al. “Diachronic Embedding for Temporal Knowledge Graph Completion”. In: *Thirty-Fourth AAAI Conference on Artificial Intelligence*. 2020, pp. 3988–3995.
- [13] Woojeong Jin et al. “Recurrent Event Network: Autoregressive Structure Inference over Temporal Knowledge Graphs”. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. 2020, pp. 6669–6683.
- [14] Timothée Lacroix, Guillaume Obozinski, and Nicolas Usunier. “Tensor Decompositions for Temporal Knowledge Base Completion”. In: *International Conference on Learning Representations*. 2020.
- [15] Ivana Balažević, Carl Allen, and Timothy M Hospedales. “Tucker: Tensor factorization for knowledge graph completion”. In: *arXiv preprint arXiv:1901.09590* (2019).

## References IV

- [16] Woojeong Jin et al. “Recurrent event network: Autoregressive structure inference over temporal knowledge graphs”. In: *arXiv preprint arXiv:1904.05530* (2019).
- [17] Alberto García-Durán, Sebastijan Dumancic, and Mathias Niepert. “Learning Sequence Encoders for Temporal Knowledge Graph Completion”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2018, pp. 4816–4821.
- [18] Julien Leblay and Melisachew Wudage Chekol. “Deriving validity time in knowledge graph”. In: *Companion Proceedings of the The Web Conference 2018*. 2018, pp. 1771–1776.
- [19] Xi Victoria Lin, Richard Socher, and Caiming Xiong. “Multi-hop knowledge graph reasoning with reward shaping”. In: *arXiv preprint arXiv:1808.10568* (2018).
- [20] Rajarshi Das et al. “Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning”. In: *arXiv preprint arXiv:1711.05851* (2017).



## References V

- [21] Sebastian Ruder. “An overview of gradient descent optimization algorithms”. In: *arXiv preprint arXiv:1609.04747* (2016).
- [22] Théo Trouillon et al. “Complex embeddings for simple link prediction”. In: *International conference on machine learning*. PMLR. 2016, pp. 2071–2080.
- [23] Elizabeth Boschee et al. *ICEWS Coded Event Data*. 2015.
- [24] Diederik P Kingma. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [25] Farzaneh Mahdisoltani, Joanna Biega, and Fabian M Suchanek. “Yago3: A knowledge base from multilingual wikipedias”. In: *CIDR*. 2013.
- [26] John Duchi, Elad Hazan, and Yoram Singer. “Adaptive subgradient methods for online learning and stochastic optimization.”. In: *Journal of machine learning research* 12.7 (2011).
- [27] Ronald J Williams. “Simple statistical gradient-following algorithms for connectionist reinforcement learning”. In: *Machine Learning* 8 (1992), pp. 229–256.