

# Evolving Beyond Snapshots: Harmonizing Structure and Sequence via Entity State Tuning for Temporal Knowledge Graph Forecasting

Siyuan Li<sup>1\*†</sup>, Yunjia Wu<sup>1†</sup>, Yiyong Xiao<sup>2†</sup>, Pingyang Huang<sup>1</sup>,  
Peize Li<sup>3</sup>, Ruitong Liu<sup>4\*</sup>, Yan Wen<sup>5</sup>, Te Sun<sup>6</sup>

<sup>1</sup> Dalian University of Technology    <sup>2</sup> Shenzhen University of Advanced Technology

<sup>3</sup> King’s College London    <sup>4</sup> Peking University

<sup>5</sup> Beijing Institute of Technology    <sup>6</sup> Shanghai Jiao Tong University

yuanlsy@mail.dlut.edu.cn, ruitong.jerry@gmail.com

## Abstract

Temporal knowledge graph (TKG) forecasting requires predicting future facts by jointly modeling structural dependencies within each snapshot and temporal evolution across snapshots. However, most existing methods are stateless: they recompute entity representations at each timestamp from a limited query window, leading to episodic amnesia and rapid decay of long-term dependencies. To address this limitation, we propose Entity State Tuning (EST), an encoder-agnostic framework that endows TKG forecasters with persistent and continuously evolving entity states. EST maintains a global state buffer and progressively aligns structural evidence with sequential signals via a closed-loop design. Specifically, a topology-aware state perceiver first injects entity-state priors into structural encoding. Then, a unified temporal context module aggregates the state-enhanced events with a pluggable sequence backbone. Subsequently, a dual-track evolution mechanism writes the updated context back to the global entity state memory, balancing plasticity against stability. Experiments on multiple benchmarks show that EST consistently improves diverse backbones and achieves state-of-the-art performance, highlighting the importance of state persistence for long-horizon TKG forecasting.

## 1 Introduction

Temporal Knowledge Graphs (TKGs) extend static facts into quadruples  $(s, r, o, t)$ , organizing knowledge as a sequence of chronological snapshots (Cai et al., 2024a; Wang et al., 2024a; Liu et al., 2025a; Lin et al., 2026). This temporal grounding is pivotal for reasoning in dynamic environments, enabling applications ranging from Question Answer-

ing (Wang et al., 2024b; Liu et al., 2026) to Recommender Systems (Wang et al., 2025).

Unlike static reasoning, TKG forecasting (Liao et al., 2024) demands extrapolating future events by deciphering the complex interplay between two orthogonal information flows: the *structural flow* (intra-snapshot topological dependencies) and the *sequential flow* (inter-snapshot temporal evolution).

Existing TKG forecasting methods typically follow a divide-and-conquer paradigm, focusing either on structural dependency modeling, exemplified by RE-GCN (Li et al., 2021), or on sequential linearization, as in TiSASRec (Li et al., 2020). Despite their architectural diversity, most of these methods still operate under a fundamentally stateless paradigm. As shown in Figure 1(a), entity representations are reconstructed independently at each snapshot, preventing historical evidence from being persistently accumulated over time. As temporal reasoning proceeds, long-range dependencies are therefore progressively weakened, leading to what we term *episodic amnesia*. By contrast, Figure 1(b) highlights the necessity of modeling entities as persistent state carriers to naturally bridge structural and temporal dynamics.

Motivated by this observation, we propose an encoder-agnostic framework that redefines TKG reasoning through Entity State Tuning (EST). We contend that the key to superior performance lies not in the choice of a specific backbone, but in effectively bridging structural and sequential flows through persistent states. Specifically, we model entities as Dynamic State Containers that persist and evolve. Our framework integrates a modular structural encoder with a pluggable sequential encoder (ranging from classic RNNs to modern Mamba) (Sherstinsky, 2020; Gu and Dao, 2024). Within this architecture, the EST mechanism acts as a semantic bridge: it continuously fuses topo-

\*Corresponding authors.

†Equal contribution.

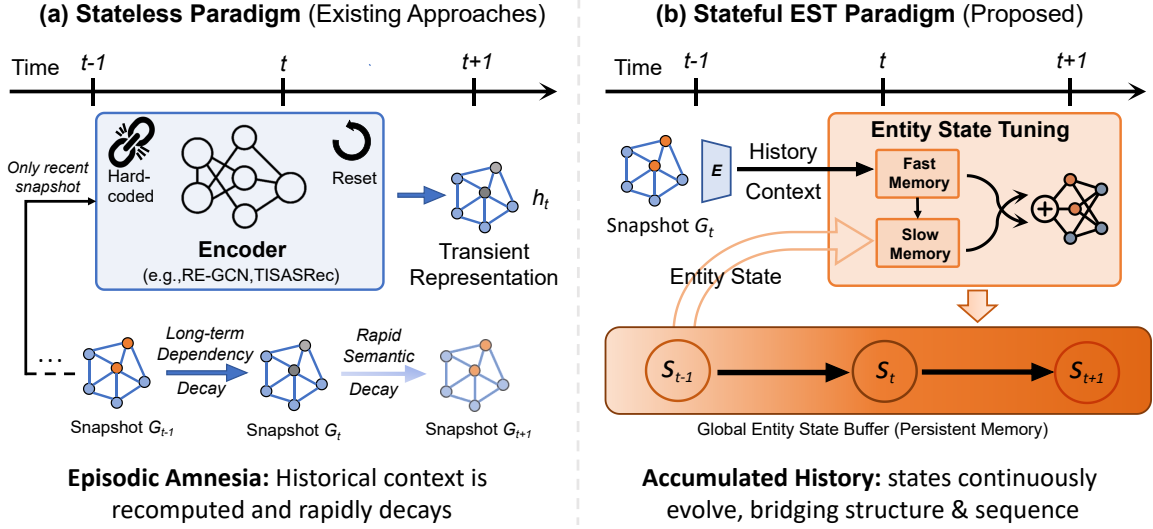


Figure 1: From Stateless Snapshot Encoding to Stateful Entity Reasoning in Temporal Knowledge Graphs.

logical context with temporal dynamics and writes the updates back into a global buffer. This ensures that entities evolve as stateful agents, preserving and accumulating historical knowledge along the timeline. Complementary to state tuning, we introduce Counterfactual Consistency Learning to correct observational visibility bias. Counterfactual constraints enable the model to isolate intrinsic temporal dependencies from spurious observational correlations. Our main contributions are summarized as follows:

- We identify the stateless bottleneck in existing TKG methods and pioneer a state-centric paradigm. This approach unifies structural and sequential dynamics by modeling entities as continuously evolving agents, preventing semantic decay over time.
- We propose Entity State Tuning, a universal framework that empowers diverse backbones with persistent memory. EST effectively harmonizes topological context with temporal evolution, boosting performance without altering the core encoder architecture.
- We introduce Counterfactual Consistency Learning to mitigate visibility bias. This objective steers the model away from spurious correlations, ensuring that the entity state evolution is grounded in temporal dependencies.
- Extensive experiments demonstrate EST’s state-of-the-art performance, empirically validating the framework’s universality and the critical necessity of state persistence for long-horizon extrapolation.

## 2 Related Works

**Temporal Knowledge Graph Forecasting.** Existing TKG forecasting literature predominantly navigates the dichotomy between *sequential evolution* and *structural reasoning* (Cai et al., 2024a; Wang et al., 2024a). Sequence-centric approaches, exemplified by RE-NET (Jin et al., 2020) and CyGNet (Zhu et al., 2021), linearize historical neighborhoods to capture repetitive temporal patterns via recurrent or copy-generation mechanisms. Conversely, structure-centric methods like RE-GCN (Li et al., 2021) and CEN (Li et al., 2022b) prioritize topological dependencies within snapshots, employing evolving GCNs to propagate information across dynamic neighbors. More recently, the field has expanded into generative and cognitive paradigms: DiffuTKG (Cai et al., 2024b) leverages diffusion processes to model future stochasticity, while emerging works in 2025, such as CognTKE (Chen et al., 2025) and TRCL (Liu et al., 2025b), incorporate cognitive dual-process theories and recursive contrastive learning to enhance extrapolation. However, a fundamental bottleneck persists: entities are treated as stateless “tabula rasa,” relying on bounded retrieval. EST transcends this by maintaining global persistent states, ensuring forecasting is driven by continuous evolutionary inertia rather than episodic context.

**Stateful Reasoning in Dynamic Graphs.** Maintaining persistent node states has been extensively explored in continuous-time dynamic graph learning. Models like JODIE (Kumar et al., 2019) and TGN (Rossi et al., 2020) utilize memory modules

to store historical interaction patterns. Recent advances like DyMemR (Zhang et al., 2024) further extend this by employing dynamic memory pools to capture repetitive historical facts. Distinct from these mechanisms, which often serve as passive storage for raw events or homogeneous interactions, EST introduces a specialized framework for stateful TKG forecasting. Unlike standard architectures that decouple memory updates from topological aggregation, EST adopts a *state-first* paradigm: it actively injects global state priors before structural encoding, creating a closed-loop evolution that mitigates semantic decay in long-term reasoning.

### 3 Problem Formulation

Let  $\mathcal{E}$  and  $\mathcal{R}$  denote the sets of entities and relations, respectively. A Temporal Knowledge Graph (TKG) is formalized as a sequence of snapshots  $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_T\}$ , where each snapshot  $\mathcal{G}_t = \{(s, r, o, t) \mid s, o \in \mathcal{E}, r \in \mathcal{R}\}$  comprises quadruples representing facts valid at timestamp  $t$ .

Given a future-time query  $q = (s, r, ?, t)$  with  $t > T_{\text{obs}}$ , the task of *TKG extrapolation* aims to predict the missing object entity by leveraging the historical interactions of the subject  $s$ . The historical context of  $s$  prior to time  $t$  is defined as  $\mathcal{H}_s(t) = \{(o_k, r_k, \tau_k) \mid (s, r_k, o_k, \tau_k) \in \mathcal{G}_{\tau_k}, \tau_k < t\}$ , characterizing the temporal evolution of relational patterns associated with the subject.

In contrast to the prevailing *stateless* paradigm which relies solely on discrete event retrieval, we postulate that each entity  $e \in \mathcal{E}$  resides on a *continuously evolving state manifold* within a latent semantic space. Let  $\mathbf{S}_t \in \mathbb{R}^{|\mathcal{E}| \times d}$  denote the global entity state matrix at time  $t$ . Our key premise is that future prediction is conditional not only on the event history  $\mathcal{H}_s(t)$  but critically on the entity’s current evolutionary momentum, encapsulated by  $\mathbf{S}_{t-1}$ . Accordingly, the learning objective reduces to maximizing the log-likelihood:

$$\max_{\Theta} \sum_{(s,r,o,t) \in \mathcal{D}} \log P_{\Theta}(o \mid s, r, \mathbf{S}_{t-1}, \mathcal{H}_s(t)). \quad (1)$$

This formulation unifies historical structural evidence with evolving latent states, enabling the model to retain long-term inertia beyond the immediate context window.

## 4 Methodology

EST is founded on the principle of *State Persistence*, comprising three tightly coupled components: (1) a Topology-Aware State Perceiver

that harmonizes local structural dependencies with global state priors; (2) a Unified Temporal Context Module that synthesizes structural and temporal signals into a continuous sequence; and (3) a Deconfounded State Evolution mechanism that closes the loop by updating global states via a dual-system memory derived from temporal guarantees.

### 4.1 Topology-Aware State Perceiver

To bridge discrete structural events with continuous entity evolution, we adopt a *State-First* encoding strategy. Unlike stateless GNNs, we posit that the interpretation of any topological dependency must be explicitly modulated by the participating entity’s current state.

**State-Augmented Entity Representation.** For a neighboring entity  $o_k$  involved in the  $k$ -th interaction, we synthesize a time-aligned representation  $\tilde{\mathbf{E}}_{o_k}$  by fusing its invariant static embedding  $\mathbf{E}_{o_k}$  with its dynamic state  $\mathbf{S}_{o_k}$  retrieved from the global buffer. An adaptive gating mechanism mitigates semantic drift by dynamically balancing intrinsic identity against evolving context:

$$\mathbf{g}_k = \sigma(\mathbf{W}_g[\mathbf{E}_{o_k} \parallel \mathbf{S}_{o_k}]), \quad (2)$$

$$\tilde{\mathbf{E}}_{o_k} = \mathbf{g}_k \odot \mathbf{E}_{o_k} + (1 - \mathbf{g}_k) \odot \mathbf{S}_{o_k}, \quad (3)$$

where  $\sigma(\cdot)$  is the sigmoid function and  $\mathbf{W}_g$  is a learnable projection.

**Structural Context Encoding.** We then model the local interaction driven by relation  $r_k$ . The state-augmented feature  $\tilde{\mathbf{E}}_{o_k}$  is fed into a structural encoder  $\Phi_{\text{struct}}$  alongside the relation embedding  $\mathbf{E}_{r_k}$  to derive the event representation:

$$\mathbf{X}_k = \Phi_{\text{struct}}(\tilde{\mathbf{E}}_{o_k}, \mathbf{E}_{r_k}). \quad (4)$$

Here,  $\Phi_{\text{struct}}$  serves as a generalized topological projector. In our implementation, we instantiate  $\Phi_{\text{struct}}$  flexibly to accommodate varying graph sparsities, ensuring the interaction is grounded in both local topology and historical state inertia.

### 4.2 Unified Temporal Context Encoder

Having derived the state-enhanced structural representations  $\{\mathbf{X}_k\}$ , we proceed to model the continuous-time dynamics of the subject’s history. This module synthesizes structural and temporal contexts into a unified sequential manifold.

**Temporal Calibration.** To account for the irregularity of event intervals, we introduce a time-difference projection operator  $\varphi(\Delta t)$ . Let  $\Delta\tau_k = t - \tau_k$  denote the elapsed time since the  $k$ -th interaction. The calibrated input  $\mathbf{u}_k$  is defined as:

$$\mathbf{u}_k = \text{MLP}([\mathbf{X}_k \parallel \mathbf{E}_{r_k} \parallel \varphi(\Delta\tau_k)]). \quad (5)$$

By explicitly re-injecting  $\mathbf{E}_{r_k}$ , we preserve relation-specific semantics that might be diluted during structural aggregation.

**Pluggable Sequence Backbone.** We formulate EST as an encoder-agnostic framework that supports diverse sequential architectures within the same state-tuning pipeline. Given the stream  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_L]$ , the sequence backbone maps each causal prefix to a contextualized representation:

$$\mathbf{h}_k = \mathcal{T}(\mathbf{u}_{\leq k}; \theta_{\text{seq}}), \quad \mathbf{y}_k = \mathcal{P}(\mathbf{h}_k), \quad (6)$$

where  $\mathbf{u}_{\leq k} = [\mathbf{u}_1, \dots, \mathbf{u}_k]$  denotes the causally visible prefix up to step  $k$ ,  $\mathbf{h}_k$  is the contextualized representation produced by the chosen backbone at position  $k$ , and  $\mathcal{P}(\cdot)$  is a projection head. For notational uniformity, we use the same form for all backbone instantiations, while their concrete formulations (e.g., RNN, LSTM, Transformer, and Mamba) are provided in **Appendix D**. We validate this flexibility empirically in Section 5.2.

**Query-Specific Attention.** To extract evidence relevant to the specific query  $(s, r, ?, t)$ , we apply a query-aware attention mechanism over the encoder output  $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_L]$ . The attention weight  $\alpha_k$  is computed as:

$$\alpha_k \propto \exp\left(\mathbf{v}^\top \tanh(\mathbf{W}_{\text{att}}[\mathbf{y}_k; \mathbf{e}_{r_k}; \varphi(\Delta\tau_k)])\right). \quad (7)$$

The final dynamic context  $\mathbf{c}_s = \sum_{k=1}^L \alpha_k \mathbf{y}_k$  aggregates historical signals weighted by their relevance to the current prediction target.

### 4.3 De-confounded State Evolution

The core innovation of EST lies in its closed-loop state evolution. To resolve the *Plasticity-Stability Dilemma* in non-stationary environments, we draw inspiration from the Complementary Learning Systems theory in neuroscience (Benna and Fusi, 2016) to design a dual-track update scheme (Kumaran et al., 2016; Kirkpatrick et al., 2017).

**Fast System (Working Memory).** We employ a high-sensitivity buffer  $\mathbf{S}_{\text{fast}}$  to capture instantaneous semantic perturbations induced by the current context  $\mathbf{c}_s$ . This transient state is updated unconditionally via an exponential moving average:

$$\mathbf{S}_{\text{fast}}^{(t)} \leftarrow (1 - \lambda) \mathbf{S}_{\text{fast}}^{(t-1)} + \lambda \mathbf{c}_s, \quad (8)$$

where  $\lambda$  controls responsiveness. This process mimics the rapid reactivity of working memory to recent stimuli.

**Slow System (Consolidated Memory).** The global persistent state  $\mathbf{S}_{\text{slow}}$  evolves only when significant structural shifts occur. We define the divergence between the fast and slow states as a measure of *informational surprise*:  $\delta_s = \|\mathbf{S}_{\text{fast}}^{(t)} - \mathbf{S}_{\text{slow}}^{(t-1)}\|_2$ . To suppress stochastic noise, we introduce an energy-barrier gating mechanism  $g_s$ :

$$g_s = \sigma(\kappa \cdot (\delta_s - \gamma)), \quad \mathbf{S}_{\text{slow}}^{(t)} \leftarrow \mathbf{S}_{\text{slow}}^{(t-1)} + g_s \cdot \left(\mathbf{S}_{\text{fast}}^{(t)} - \mathbf{S}_{\text{slow}}^{(t-1)}\right), \quad (9)$$

where  $\gamma$  is the cognitive threshold. This ensures the entity manifold is reconfigured only when accumulated evidence surpasses the energy barrier ( $\delta_s > \gamma$ ), distilling robust knowledge while filtering transient noise.

**Counterfactual Consistency Learning.** To prevent the model from capturing spurious correlations (e.g., merely memorizing high-frequency historical entities), we employ a *Counterfactual Negative Sampling* strategy to construct the negative set  $\mathcal{E}^-$ . Specifically, for a valid fact  $(s, r, o, t)$ , we sample negative entities  $o'$  that are historically plausible (e.g., have interacted with  $s$  in the past) but are factually incorrect at the current timestamp  $t$ . The optimization is to minimize the following negative log-likelihood loss:

$$\mathcal{L}_{\text{ccl}} = - \mathbb{E}_{(s,r,o,t) \sim \mathcal{D}} \left[ \phi(\mathbf{c}_s, r, o) - \log \sum_{o' \in \{o\} \cup \mathcal{E}^-} \exp(\phi(\mathbf{c}_s, r, o')) \right], \quad (10)$$

where  $\phi(\cdot)$  denotes the scoring function (variants detailed in **Appendix C**). This normalization over  $\{o\} \cup \mathcal{E}^-$  compels the model to distinguish intrinsic temporal dependencies from historical priors.

## 5 Experiments

We empirically validate EST by addressing the following research questions:

Table 1: Overall performance on four benchmark datasets. Best results are in **bold**, second best are underlined.

Model	ICEWS14				ICEWS18				ICEWS05-15				GDELT				
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10	
RE-NET (2020)	0.369	0.268	0.395	0.548	0.288	0.191	0.324	0.475	0.433	0.334	0.478	0.631	0.196	0.124	0.210	0.340	
CyGNet (2020)	0.351	0.257	0.390	0.536	0.249	0.159	0.283	0.426	0.368	0.266	0.416	0.562	0.185	0.115	0.196	0.320	
xERTE (2021)	0.400	0.321	0.446	0.562	0.293	0.210	0.335	0.465	0.466	0.378	0.523	0.639	0.181	0.123	0.201	0.303	
TITer (2021)	0.409	0.323	0.455	0.571	0.300	0.221	0.335	0.448	0.477	0.380	0.529	0.658	0.155	0.110	0.156	0.243	
RE-GCN (2021)	0.404	0.307	0.450	0.592	0.306	0.210	0.343	0.488	0.480	0.373	0.539	0.683	0.196	0.124	0.209	0.337	
CEN (2022)	0.422	0.321	0.475	0.613	0.315	0.217	0.354	0.506	0.468	0.364	0.524	0.670	0.204	0.130	0.218	0.350	
TiRGN (2022)	0.440	0.338	0.490	0.638	0.337	0.232	0.380	0.542	0.500	0.393	0.561	0.707	0.217	0.136	0.233	0.376	
HisMatch (2022)	0.464	0.359	0.516	0.668	0.340	0.239	0.379	0.539	0.528	0.420	0.591	0.733	0.220	0.145	0.238	0.366	
RETIA (2023)	0.428	0.323	0.478	0.628	0.324	0.222	0.365	0.529	0.473	0.366	0.529	0.678	0.201	0.128	0.215	0.345	
CENET (2023)	0.390	0.296	0.432	0.575	0.279	0.182	0.316	0.470	0.420	0.322	0.469	0.604	0.202	0.127	0.217	0.349	
PPT (2023)	0.384	0.289	0.425	0.570	0.266	0.169	0.306	0.454	0.388	0.286	0.434	0.586	–	–	–	–	
LogCL (2023)	<u>0.489</u>	<u>0.378</u>	<b>0.547</b>	<b>0.703</b>	0.357	0.245	0.403	0.577	0.570	0.461	<u>0.637</u>	<b>0.779</b>	0.238	0.146	0.256	0.423	
Re-Temp (2023)	0.480	0.373	0.536	<u>0.689</u>	0.358	0.250	0.404	0.573	0.563	0.455	0.628	<u>0.772</u>	0.250	0.157	0.271	0.442	
DiffuTKG (2024)	–	–	–	–	0.367	0.257	–	0.578	–	0.527	0.404	–	0.760	0.251	0.162	–	0.423
CognTKE (2025)	0.461	0.365	0.511	0.645	0.352	0.252	0.399	0.547	0.531	0.426	0.594	0.727	–	–	–	–	
DSEP (2025)	0.449	0.344	0.503	0.649	0.338	0.232	0.382	0.545	0.507	0.398	0.568	0.713	0.220	0.139	0.237	0.381	
EST-RNN	0.501	0.447	0.524	0.618	0.458	0.398	<u>0.480</u>	0.575	<u>0.611</u>	<u>0.560</u>	0.631	0.712	<u>0.408</u>	0.348	<u>0.426</u>	<u>0.510</u>	
EST-LSTM	0.505	0.448	0.532	0.625	<b>0.467</b>	<b>0.407</b>	<b>0.490</b>	<b>0.585</b>	0.609	0.558	0.629	0.710	<u>0.408</u>	<u>0.350</u>	0.425	<b>0.519</b>	
EST-Transformer	<b>0.513</b>	<b>0.455</b>	0.536	0.629	<u>0.466</u>	<u>0.405</u>	<b>0.490</b>	<u>0.583</u>	<b>0.621</b>	<b>0.572</b>	<b>0.640</b>	0.718	<b>0.412</b>	<b>0.356</b>	<b>0.429</b>	<b>0.519</b>	
EST-Mamba	<u>0.511</u>	<u>0.451</u>	<u>0.538</u>	0.633	0.449	0.383	0.475	0.576	0.599	0.544	0.624	0.709	0.390	0.329	0.408	0.504	

- **RQ1:** How does EST perform compared to representative TKG models?
- **RQ2:** What are the contributions of EST’s key components, and how sensitive is the model to key hyperparameters?
- **RQ3:** How does the EST model the evolution of entity states over time, and how do structural and sequential signals interact within this process?
- **RQ4:** Can the learned entity states generalize across time beyond the training distribution?

## 5.1 Experiment Settings

**Datasets.** We evaluate EST on four standard benchmarks—**ICEWS14**, **ICEWS18**, **ICEWS05-15**, and **GDELT**—spanning diverse scales of political and global event dynamics. Detailed dataset statistics and implementation details are provided in **Appendix A and B**.

**Baselines.** We compare against a broad set of TKG models, including RE-NET(Jin et al., 2020), CyGNet(Zhu et al., 2021), xERTE(Han et al., 2020), TITer(Sun et al., 2021), RE-GCN(Li et al., 2021), CEN(Li et al., 2022b), TiRGN(Li et al., 2022a), HiSMatch(Li et al., 2022c), RETIA(Liu et al., 2023), CENET(Xu et al., 2023b), PPT(Xu et al., 2023a), LogCL(Chen et al., 2024), Re-Temp(Wang et al., 2023), DiffuTKG(Cai et al., 2024b), CognTKE(Chen et al., 2025), and DSEP(Deng et al., 2025).

**Evaluation Metrics.** We adopt the standard filtered MRR and Hits@N (H@N) for evaluation.

## 5.2 Overall Performance Comparison: RQ1


**Main results.** Table 1 demonstrates that the proposed EST framework consistently establishes new state-of-the-art results across all benchmarks, with EST-Transformer achieving an MRR of 0.513 on ICEWS14 and a remarkable 0.412 on GDELT. Most notably, on the high-noise GDELT dataset, our model yields a relative improvement of over **60%** compared to leading baselines (e.g., DiffuTKG, Re-Temp), strongly validating that our *persistent state mechanism* effectively mitigates the “episodic amnesia” inherent in stateless methods by robustly filtering long-term noise. Crucially, despite utilizing a lightweight MLP for structural encoding ( $\Phi_{\text{struct}}$ ), even the basic EST-RNN variant outperforms structurally complex models (e.g., TiRGN, DSEP). This underscores that the bottleneck of TKG forecasting lies in the reasoning paradigm rather than encoder complexity. The consistent superiority across diverse backbones (from RNN to Mamba) further confirms EST as a highly adaptable framework for temporal reasoning.

### 5.2.1 Efficiency and Scalability


EST is encoder-agnostic, abstracting temporal evolution as a generalized state transition function  $\mathcal{T}$ . We instantiate  $\mathcal{T}$  with four backbones; the corresponding formulations are provided in **Appendix D**. As shown in Figure 2 and Table 2, all variants converge reliably, indicating that the gains primarily come from persistent state modeling rather than a particular sequence encoder. Among them, EST-Transformer achieves the best overall accuracy, while EST-Mamba offers the most

Table 2: Cumulative wall-clock training time (minutes) at different training epochs on the ICEWS14 dataset.

Model	Training Epochs													
	Start	1	2	3	4	5	6	7	8	9	10	15	20	
Ours w/ RNN	0.00	1.07	2.12	3.17	4.18	5.22	6.27	7.32	8.35	9.38	10.42	15.62	20.85	
Ours w/ LSTM	0.00	1.13	2.22	3.32	4.40	5.48	6.60	7.70	8.78	9.88	11.00	16.52	22.03	
Ours w/ Transformer	0.00	0.95	1.90	2.82	3.75	4.68	5.62	6.55	7.48	8.42	9.35	13.92	18.58	
Ours w/ Mamba	0.00	1.03	2.02	3.01	4.00	5.00	5.98	6.97	7.98	8.95	9.95	14.90	19.81	



(a) Test MRR convergence.



(b) Model size comparison (MB).

Figure 2: Efficiency comparison among different sequence encoders on the ICEWS14 dataset.

Table 3: Additional memory introduced by the two FP32 global state buffers in EST ( $d=64$ ).

Dataset	#Entities	Buffer memory (FP32)
ICEWS14	7,128	3.48 MB
ICEWS18	10,094	4.93 MB
ICEWS05-15	23,033	11.25 MB
GDELT	7,691	3.76 MB
Extrapolation	1,000,000	488.28 MB (~0.49 GB)

Table 4: Measured runtime and peak-memory overhead of EST on ICEWS14, compared with a matched stateless ablation (**w/o State**).

Setting	s/epoch	ms/step	Peak MB
EST-Mamba	60.77	105.09	173.80
w/o State	59.60	102.46	165.21
$\Delta$	+1.96%	+2.57%	+5.20%
EST-Transformer	61.51	105.22	111.61
w/o State	61.04	104.96	101.12
$\Delta$	+0.77%	+0.25%	+10.37%

parameter-efficient alternative with competitive performance. Accordingly, we use EST-Transformer as the default accuracy-oriented instantiation and EST-Mamba as a lightweight counterpart in subsequent analyses.

Beyond encoder-level efficiency, we further analyze the overhead introduced by persistent global states. EST maintains two *non-trainable* entity buffers,  $\mathbf{S}_{\text{fast}}$  and  $\mathbf{S}_{\text{slow}}$ , each of size  $|\mathcal{E}| \times d$ . The resulting memory overhead is therefore linear in the number of entities, i.e.,  $\mathcal{O}(|\mathcal{E}|d)$ . Under our default setting ( $d=64$ , FP32), the total buffer footprint is

$2 \times |\mathcal{E}| \times 64 \times 4 = 512|\mathcal{E}|$  bytes, which remains at MB scale on standard benchmarks (Table 3).

Importantly, EST does not scan all entity states at each step. State access is sparse: for each batch, we only read and update entities appearing in the current history window. As a result, the runtime remains dominated by the sequential backbone and local structural encoding, while the global buffers introduce only lightweight indexing and in-place updates. Accordingly, the overall sequence modeling complexity remains that of the chosen backbone, namely  $\mathcal{O}(Ld^2)$  for RNN/LSTM,  $\mathcal{O}(L^2d)$  for Transformer, and  $\mathcal{O}(Ld)$  for Mamba.

To quantify the practical overhead of persistent states, we compare EST against a strictly matched stateless ablation (**w/o State**) on ICEWS14. Table 4 shows that the additional runtime is modest for both backbones: for Mamba, EST introduces only +1.96% s/epoch and +2.57% ms/step, while for Transformer the increase is +0.77% and +0.25%, respectively. The peak-memory increase is also limited (+5.20% for Mamba and +10.37% for Transformer). Overall, EST introduces a linear memory overhead in  $|\mathcal{E}|$  and only minor runtime cost in practice, making the global-buffer design practical on standard TKGs and viable at larger scales.

### 5.3 Ablation Study: RQ2

To verify the effectiveness of our EST, we develop three distinct model variants for ablation studies:

Table 5: Ablation study of key components on the ICEWS14 and ICEWS18 datasets.

Model	MRR	Hits@1	Hits@3	Hits@10
<b>ICEWS14</b>				
EST-Mamba	<b>0.511</b>	<b>0.451</b>	<b>0.538</b>	<b>0.633</b>
w/o Context	0.360	0.266	0.403	0.536
w/o State	0.452	0.384	0.481	0.587
w/o CCL	0.489	0.427	0.511	0.613
<b>ICEWS18</b>				
EST-Mamba	<b>0.449</b>	<b>0.383</b>	<b>0.475</b>	<b>0.576</b>
w/o Context	0.259	0.165	0.295	0.442
w/o State	0.381	0.308	0.412	0.521
w/o CCL	0.433	0.380	0.468	0.557

Model	MRR	Hits@1	Hits@3	Hits@10
<b>ICEWS14</b>				
EST-Transformer	<b>0.513</b>	<b>0.455</b>	<b>0.536</b>	<b>0.629</b>
w/o Context	0.360	0.266	0.403	0.536
w/o State	0.455	0.386	0.479	0.582
w/o CCL	0.502	0.445	0.531	0.615
<b>ICEWS18</b>				
EST-Transformer	<b>0.466</b>	<b>0.405</b>	<b>0.490</b>	<b>0.583</b>
w/o Context	0.259	0.165	0.295	0.442
w/o State	0.397	0.326	0.424	0.529
w/o CCL	0.449	0.396	0.487	0.559

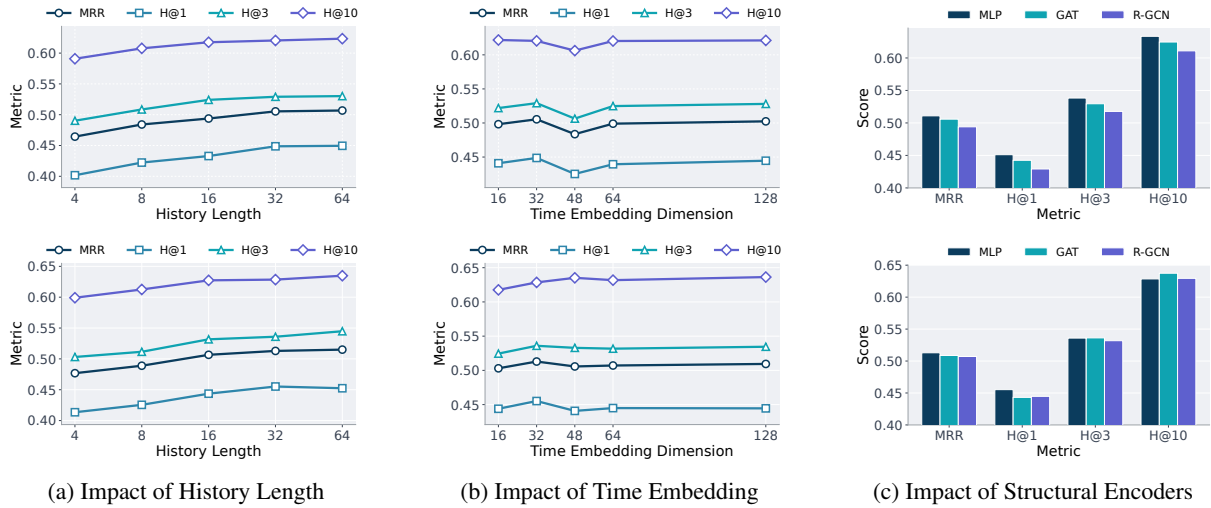


Figure 3: Hyperparameter analysis for EST-Mamba (top) and EST-Transformer (bottom).

- **w/o Context:** We remove both structural and sequential encoders, predicting from the query  $(s, r, ?, t)$  alone.
- **w/o State:** We keep the context encoders but disable state retrieval and closed-loop state updates, using only static embeddings.
- **w/o CCL:** We remove the counterfactual consistency objective and train with the standard likelihood loss using conventional negative sampling.

### 5.3.1 Key Modules Ablation Study

Table 5 dissects the impact of individual components across ICEWS14 and ICEWS18. We observe a clear hierarchical dependency in component efficacy. Removing the historical context (**w/o Context**) causes a precipitous performance drop (e.g., MRR declines from 0.511 to 0.360 on ICEWS14), confirming the foundational role of historical evidence over static query matching. Crucially, decoupling the state persistence mechanism (**w/o State**) leads to a substantial degradation (average  $\Delta$ MRR  $\approx -0.06$ ) across all settings. This empirical gap directly validates our core hypothesis: without the

persistent state to bridge temporal disjoints, the model suffers from rapid semantic decay. Furthermore, the removal of counterfactual consistency (**w/o CCL**) results in a consistent performance dip (up to 0.022 MRR), underscoring its necessity as a regularizer to filter spurious historical priors. These trends persist across Mamba and Transformer backbones, indicating EST’s backbone-agnostic design.

### 5.3.2 Sensitivity to Key Hyperparameters

Figure 3 studies the effects of history length  $L$ , time-embedding dimension, and structural encoders on ICEWS14. Overall, both variants show consistent trends, indicating that EST is stable to these design choices.

**History length.** Increasing  $L$  steadily improves MRR and Hits, with gains saturating around  $L = 32$ ; longer histories bring marginal benefits. This suggests EST can effectively exploit longer context, while remaining robust under short windows.

**Time embedding dimension.** Performance is largely insensitive to the time-embedding size. Both models peak around moderate dimensions

Table 6: Top-5 entities exhibiting the largest state displacement on the ICEWS14 dataset.

Rank	Entity	Gate $\mu$	Gate $\sigma$
1	Men_(Namibia)	0.8875	0.1507
2	Political_Parties_(Indonesia)	0.8742	0.1737
3	Police_(Kosovo)	0.8865	0.1435
4	Criminal_(Brazil)	0.8702	0.1800
5	Will_Hodgman	0.9126	0.1487

Table 7: Top-5 entities exhibiting the largest state displacement on the ICEWS18 dataset.

Rank	Entity	Gate $\mu$	Gate $\sigma$
1	Police (Honduras)	0.8740	0.1653
2	Protester (Honduras)	0.8638	0.1767
3	Police (Malawi)	0.8698	0.1549
4	Protester (Tunisia)	0.8837	0.1573
5	Citizen (Malawi)	0.8721	0.1762

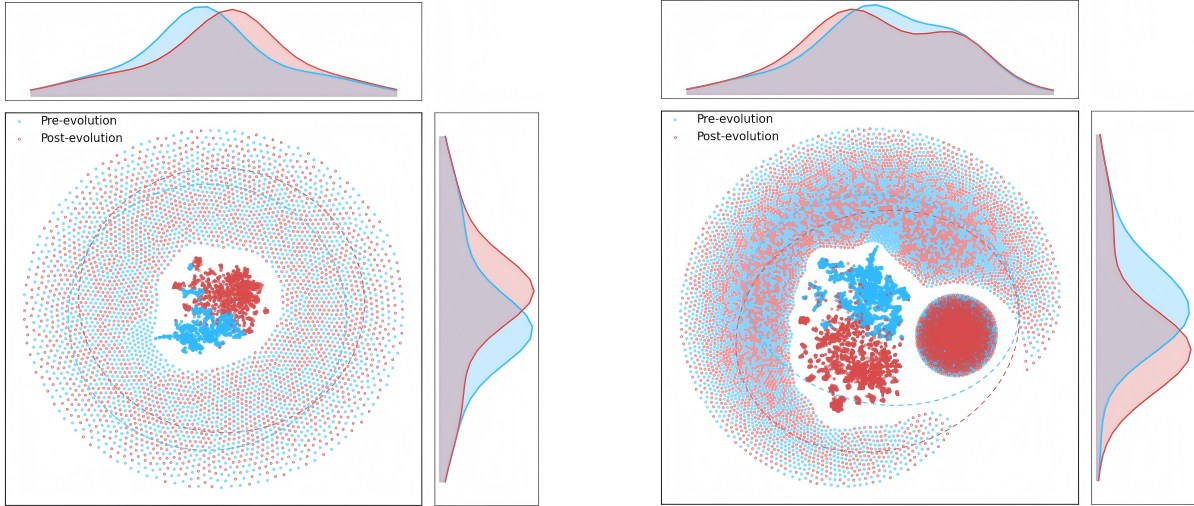


Figure 4: t-SNE visualizations of entity states before and after evolution on ICEWS14 (left) and ICEWS18 (right).

(e.g., 32) and remain stable thereafter, with only minor fluctuations at intermediate sizes.

**Structural encoders.** Swapping the structural module among MLP, GAT, and R-GCN yields only small differences. A lightweight MLP is already competitive (often best on MRR), implying that the dominant improvements are driven by EST’s state persistence and temporal modeling rather than a heavily engineered structural encoder.

#### 5.4 Entity State Evolution Analysis: RQ3

To visualize how entity states evolve over time, we track the evolution trajectories of a random 50% subset of entities and project their persistent states with t-SNE (Maaten and Hinton, 2008). As shown in Figure 4, pre- and post-evolution states largely overlap for most entities, while a relatively small portion exhibits clear displacement and local cluster reconfiguration. This pattern indicates that EST preserves a stable global geometry and selectively adjusts states when new evidence accumulates.


Motivated by the above, we quantify *state displacement* as the  $\ell_2$  change of an entity’s persistent state across updates and extract the top movers. Tables 6–7 list the top-5 entities with the largest displacement on ICEWS14/18, which

correspond to highly eventful actors (e.g., *Police/Protester/Citizen* in specific countries). For these entities, we further report the statistics (Gate  $\mu/\sigma$ ) of the *state-augmentation gate*  $\mathbf{g}_k$  in Eq. (2), which controls how static identity embeddings and dynamic states are fused when constructing state-aware neighborhood representations. The consistently high Gate  $\mu$  with moderate variance suggests that, even for rapidly changing entities, EST keeps representations anchored to intrinsic identity while injecting state information in a controlled manner, mitigating semantic drift. Overall, the t-SNE geometry and top-displacement analysis jointly illustrate a closed interaction loop: evolved states shape future structural encoding, which in turn drives subsequent state updates.


#### 5.5 Temporal Generalization: RQ4

To evaluate robustness under temporal distribution shift, we adopt a *Chronological Data Truncation* protocol, where models are trained only on the earliest  $X\%$  of the training history ( $X \in \{50, 60, 70, 80, 90\}$ ).


Figure 5 shows that EST consistently outperforms its stateless counterpart across all truncation levels on both ICEWS14 and ICEWS18, for both




(a) EST-Transformer on ICEWS14.



(b) EST-Transformer on ICEWS18.



(c) EST-Mamba on ICEWS14.



(d) EST-Mamba on ICEWS18.

Figure 5: Temporal generalization under chronological data truncation. The first row shows EST-Transformer and the second row shows EST-Mamba; the left and right columns correspond to ICEWS14 and ICEWS18, respectively. EST consistently maintains a clear advantage over its stateless ablation across different training-history budgets.

Transformer and Mamba backbones. Although performance degrades as less history is observed, the advantage of persistent states remains stable, indicating that the benefit of EST is not tied to a particular sequence encoder but is intrinsic to the state-centric paradigm.

Notably, EST also exhibits strong data efficiency. On ICEWS18 with only 50% of the training history, EST-Transformer still achieves an MRR of 0.416, surpassing fully-supervised baselines such as DiffuTKG (0.367), while EST-Mamba reaches 0.401, also exceeding DiffuTKG and CognTKE (0.352). These results suggest that persistent entity states capture transferable evolutionary dynamics beyond the observed window, leading to stronger temporal generalization than stateless modeling.

## 6 Conclusion

We identify a fundamental limitation in temporal knowledge graph forecasting: entities are commonly modeled as stateless across time, leading to repeated reconstruction and poor long-horizon generalization. To address this, we propose Entity State Tuning (EST), an encoder-agnostic framework that equips forecasters with persistent and continuously updated entity states. EST injects state priors into structural encoding, aggregates state-aware events with a pluggable temporal backbone, and updates entity states through an evolution

mechanism that balances stability and plasticity, guided by a counterfactual objective. Experiments on four benchmarks show that EST consistently improves diverse backbones and achieves state-of-the-art results. Our findings highlight state persistence as central to robust TKG forecasting and suggest richer state evolution and memory consolidation beyond snapshot-based reasoning.

## Limitations

Notwithstanding EST’s design as a universal and efficient paradigm, capturing the full spectrum of complex, non-stationary temporal dynamics remains a non-trivial challenge

First, it maintains global entity state buffers, introducing a memory overhead that scales linearly with the number of entities, i.e.,  $\mathcal{O}(|\mathcal{E}|d)$ ; while practical on standard benchmarks, this may become costly for web-scale graphs or memory-constrained deployments. Second, EST assumes temporally ordered processing and a relatively stable entity universe, and thus does not explicitly address previously unseen entities or highly dynamic entity sets. Finally, our evaluation is limited to four standard benchmarks under chronological splits; real-world TKGs may exhibit stronger non-stationarity, exogenous shocks, or richer supervision signals (e.g., attributes and text), which are beyond the scope of the current study.

## References

- Marcus K Benna and Stefano Fusi. 2016. Computational principles of synaptic memory consolidation. *Nature neuroscience*, 19(12):1697–1706.
- Borui Cai, Yong Xiang, Longxiang Gao, He Zhang, Yunfeng Li, and Jianxin Li. 2024a. Temporal knowledge graph completion: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 36(2):772–794.
- Yuxiang Cai, Qiao Liu, Yanglei Gan, Changlin Li, Xueyi Liu, Run Lin, Da Luo, and JiayeYang JiayeYang. 2024b. Predicting the unpredictable: Uncertainty-aware reasoning over temporal knowledge graphs via diffusion process. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 5766–5778.
- Wei Chen, Huaiyu Wan, Yuting Wu, Shuyuan Zhao, Jiayaqi Cheng, Yuxin Li, and Youfang Lin. 2024. Local-global history-aware contrastive learning for temporal knowledge graph reasoning. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pages 733–746. IEEE.
- Wei Chen, Yuting Wu, Shuhan Wu, Zhiyu Zhang, Mengqi Liao, Youfang Lin, and Huaiyu Wan. 2025. Cogntke: A cognitive temporal knowledge extrapolation framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 14815–14823.
- Yunteng Deng, Jia Song, Zhongliang Yang, Yilin Long, Li Zeng, and Linna Zhou. 2025. Diachronic semantic encoding based on pre-trained language model for temporal knowledge graph reasoning. *Knowledge-Based Systems*, 318:113479.
- Alberto García-Durán, Sebastijan Dumančić, and Mathias Niepert. 2018. Learning sequence encoders for temporal knowledge graph completion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2071–2080.
- Albert Gu and Tri Dao. 2024. Mamba: Linear-time sequence modeling with selective state spaces. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235, pages 16728–16746. PMLR.
- Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. 2020. Explainable subgraph reasoning for forecasting on temporal knowledge graphs. In *International conference on learning representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Woojeong Jin, Meng Qu, Xisen Jin, and Xiang Ren. 2020. Recurrent event network: Autoregressive structure inference over temporal knowledge graphs. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 6669–6683.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, and 1 others. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Srijan Kumar, Xikun Zhang, and Jure Leskovec. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1269–1278.
- Dharshan Kumaran, Demis Hassabis, and James L McClelland. 2016. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7):512–534.
- Kalev Leetaru and Philip A Schrod. 2013. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA Annual Convention*, volume 2, pages 1–49. International Studies Association San Francisco.
- Jiacheng Li, Yujie Wang, and Julian McAuley. 2020. Time interval aware self-attention for sequential recommendation. In *Proceedings of the 13th international conference on web search and data mining*, pages 322–330.
- Yujia Li, Shiliang Sun, and Jing Zhao. 2022a. Tirgn: Time-guided recurrent graph network with local-global historical patterns for temporal knowledge graph reasoning. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)*, pages 2152–2158. International Joint Conferences on Artificial Intelligence Organization.
- Zixuan Li, Saiping Guan, Xiaolong Jin, Weihua Peng, Yajuan Lyu, Yong Zhu, Long Bai, Wei Li, and 1 others. 2022b. Complex evolutionary pattern learning for temporal knowledge graph reasoning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 6516–6526.
- Zixuan Li, Zhongni Hou, Saiping Guan, Xiaolong Jin, Weihua Peng, Long Bai, Yajuan Lyu, Wei Li, and 1 others. 2022c. Hismatch: Historical structure matching based temporal knowledge graph reasoning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2839–2850.
- Zixuan Li, Xiaolong Jin, Wei Li, Saiping Guan, Jiafeng Guo, Huawei Shen, Yuanzhuo Wang, and Xueqi Cheng. 2021. Temporal knowledge graph reasoning based on evolutionary representation learning. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 408–417.
- Ruotong Liao, Xu Jia, Yangzhe Li, Yunpu Ma, and Volker Tresp. 2024. Gentkg: Generative forecasting

- on temporal knowledge graph with large language models. In *Findings of the association for computational linguistics: NAACL 2024*, pages 4303–4317.
- Dongying Lin, Yanan Liu, Bin Wang, Xiaochun Yang, and 1 others. 2026. Ontotkge: Ontology-enhanced temporal knowledge graph extrapolation. *arXiv preprint arXiv:2604.05468*.
- Kangzheng Liu, Feng Zhao, Guandong Xu, Xianzhi Wang, and Hai Jin. 2023. Retia: relation-entity twin-interact aggregation for temporal knowledge graph extrapolation. In *2023 IEEE 39th international conference on data engineering (ICDE)*, pages 1761–1774. IEEE.
- Lihui Liu, Zihao Wang, and Hanghang Tong. 2025a. Neural-symbolic reasoning over knowledge graphs: A survey from a query perspective. *ACM SIGKDD Explorations Newsletter*, 27(1):124–136.
- Weitong Liu, Khairunnisa Hasikin, Anis Salwa Mohd Khairuddin, Meizhen Liu, and Xuechen Zhao. 2025b. A temporal knowledge graph reasoning model based on recurrent encoding and contrastive learning. *PeerJ Computer Science*, 11:e2595.
- Yinan Liu, Dongying Lin, Sigang Luo, Xiaochun Yang, and Bin Wang. 2026. Joint knowledge base completion and question answering by combining large language models and small language models. *arXiv preprint arXiv:2604.05875*.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. 2020. Temporal graph networks for deep learning on dynamic graphs. In *ICML 2020 Workshop on Graph Representation Learning*.
- Alex Sherstinsky. 2020. Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena*, 404:132306.
- Haohai Sun, Jialun Zhong, Yunpu Ma, Zhen Han, and Kun He. 2021. Timetraveler: Reinforcement learning for temporal knowledge graph forecasting. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 11065–11078.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Jiapu Wang, Boyue Wang, Meikang Qiu, Shirui Pan, Bo Xiong, Heng Liu, Linhao Luo, and 1 others. 2024a. A survey on temporal knowledge graph completion: Taxonomy, progress, and prospects. *IEEE Transactions on Knowledge and Data Engineering*, 36(10):5687–5707.
- Kunze Wang, Caren Han, and Josiah Poon. 2023. Re-temp: Relation-aware temporal representation learning for temporal knowledge graph completion. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 258–269.
- Shijie Wang, Wenqi Fan, Yue Feng, Lin Shanru, Xinyu Ma, Shuaiqiang Wang, and Dawei Yin. 2025. Knowledge graph retrieval-augmented generation for llm-based recommendation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL), ACL 2025 Long Papers*, pages 27152–27168. Association for Computational Linguistics.
- Yu Wang, Nedim Lipka, Ryan A Rossi, Alexa Siu, Ruiyi Zhang, and Tyler Derr. 2024b. Knowledge graph prompting for multi-document question answering. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, pages 19206–19214. AAAI Press.
- Wenjie Xu, Ben Liu, Miao Peng, Xu Jia, and Min Peng. 2023a. Pre-trained language model with prompts for temporal knowledge graph completion. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11041–11054.
- Yi Xu, Junjie Ou, Hui Xu, and Luoyi Fu. 2023b. Temporal knowledge graph reasoning with historical contrastive learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 4765–4773.
- Fuwei Zhang, Zhao Zhang, Fuzhen Zhuang, Yu Zhao, Deqing Wang, and Hongwei Zheng. 2024. Temporal knowledge graph reasoning with dynamic memory enhancement. *IEEE Transactions on Knowledge and Data Engineering*, 36(11):7115–7128.
- Cunchao Zhu, Muhao Chen, Changjun Fan, Guangquan Cheng, and Yan Zhang. 2021. Learning from history: Modeling temporal knowledge graphs with sequential copy-generation networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 4732–4740.

## Appendix

<b>A Datasets</b>	<b>12</b>
<b>B Implementation Details</b>	<b>12</b>
B.1 Experiment Settings . . . . .	12
B.2 Optimization. . . . .	12
B.3 Evaluation protocol. . . . .	12
<b>C Scoring Function</b>	<b>12</b>
<b>D Sequential Backbones</b>	<b>13</b>
<b>E End-to-End Workflow on ICEWS14</b>	<b>13</b>

### A Datasets

We evaluate our approach on four widely used real-world temporal knowledge graph benchmarks: ICEWS14 (García-Durán et al., 2018), ICEWS18 (Jin et al., 2020), ICEWS05-15 (Li et al., 2021), and GDELT (Leetaru and Schrod, 2013). The three ICEWS datasets—ICEWS14, ICEWS18, and ICEWS05-15—are derived from the Integrated Crisis Early Warning System and consist of large-scale international political events spanning different temporal ranges. In contrast, GDELT is constructed from global news media and covers a broader spectrum of societal events across countries and regions.

For all datasets, we adopt the standard chronological splitting protocol with an 8:1:1 ratio for training, validation, and testing (Jin et al., 2020). To prevent temporal leakage, validation data are strictly drawn from timestamps following the training period, and test data are restricted to timestamps occurring after the validation window. Detailed dataset statistics, including the numbers of entities, relations, timestamps, and quadruples in each split, are reported in Table 8.

### B Implementation Details

#### B.1 Experiment Settings

Unless otherwise stated, all models are implemented in PyTorch and trained on a single NVIDIA GeForce RTX 4090 GPU. We set the embedding dimension to  $d = 64$  and the time encoding dimension to  $d_t = 32$ . Notably, we employ a lightweight MLP as the default structural encoder  $\Phi_{\text{struct}}$ , ensuring that performance gains are attributed to the state tuning mechanism rather than complex topological aggregation. For each query at time  $\tau$ , we retrieve an entity-centric history of length  $L = 32$  (strictly

earlier than  $\tau$ ). Temporal gaps are encoded via a  $\log(1 + \Delta t)$  projection, and DistMult is adopted as the default scoring function.

To realize the proposed dual-system memory, we maintain two non-trainable entity state buffers, both of which are initialized to all-zero vectors for every entity at the beginning of training. During training, the fast state is updated using an exponential moving average with update rate  $\lambda = 0.2$ , while the slow state is updated via a sigmoid-gated interpolation with fixed scale  $\kappa = 5.0$  and threshold  $\gamma = 0.5$ . Overall, the end-to-end training and the closed-loop state evolution strictly follow **Algorithm 1**.

#### B.2 Optimization.

We train for 20 epochs with mini-batch size 128 using Adam with learning rate  $5 \times 10^{-3}$  and weight decay  $10^{-4}$ . We adopt a linear warmup for the first 2 epochs and then apply cosine annealing scheduling, where the minimum learning rate is set to 0.1 times the base learning rate.

#### B.3 Evaluation protocol.

We report Mean Reciprocal Rank (MRR) and Hits@K ( $K \in \{1, 3, 10\}$ ) under a time-consistent filtered ranking protocol, analogous to the standard filtered evaluation setting commonly adopted in TKG. Specifically, we evaluate queries in chronological order and only filter candidate tails that were observed strictly before the current timestamp (rolling setting), then reveal events at the current timestamp after evaluation. This avoids temporal leakage in filtered ranking and ensures counterfactual consistency.

### C Scoring Function

We instantiate the decoder with four scoring functions (DistMult, RotatE, ComplEx, and an MLP). On ICEWS14, DistMult consistently achieves the best results for both EST-Mamba (MRR 0.511, Hits@1 0.451, Hits@3 0.538, Hits@10 0.633) and EST-Transformer (MRR 0.513, Hits@1 0.455, Hits@3 0.536, Hits@10 0.629). This suggests that the state-enhanced temporal context produced by EST can be effectively decoded by a lightweight bilinear scorer, while more expressive complex/rotational decoders do not yield additional benefits. We therefore adopt DistMult as the default scorer in all experiments.

Table 8: Statistics of Four Temporal Knowledge Graph Benchmarks

Dataset	#Entities	#Relations	#Train	#Valid	#Test	#Snapshots
ICEWS14	7,128	230	74,845	8,514	7,371	365
ICEWS18	10,094	256	373,018	45,995	49,545	365
ICEWS05-15	23,033	251	368,868	46,302	46,159	4,017
GDELT	7,691	240	1,734,399	238,765	305,241	2,975

Table 9: Impact of scoring functions on ICEWS14.

(a) EST-Mamba on ICEWS14.

Scorer	MRR	Hit@1	Hit@3	Hit@10
DistMult	<b>0.511</b>	<b>0.451</b>	<b>0.538</b>	<b>0.633</b>
MLP	<u>0.485</u>	<u>0.416</u>	<u>0.516</u>	<u>0.618</u>
ComplEx	0.439	0.371	0.466	0.570
RotatE	0.471	0.399	0.504	0.609

(b) EST-Transformer on ICEWS14.

Scorer	MRR	Hit@1	Hit@3	Hit@10
DistMult	<b>0.513</b>	<b>0.455</b>	<b>0.536</b>	<b>0.629</b>
MLP	<u>0.475</u>	<u>0.408</u>	<u>0.502</u>	<u>0.607</u>
ComplEx	0.434	0.366	0.460	0.566
RotatE	0.470	0.400	0.501	0.604

## D Sequential Backbones

As outlined in Section 4.2, the EST framework is designed to be encoder-agnostic, abstracting the temporal evolution into a generalized state transition function  $\mathcal{T}$ . Formally, given the sequence of structure-aware event representations  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_L] \in \mathbb{R}^{L \times d}$  derived from the Unified Temporal Context Module, the sequential encoder maps  $\mathbf{U}$  to a sequence of contextualized hidden states  $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_L]$  under the constraint of temporal visibility (i.e.,  $\mathbf{h}_t$  depends only on  $\mathbf{u}_{\leq t}$ ).

In this section, we provide the specific mathematical formulations for the four instantiated backbones: RNN (Sherstinsky, 2020), LSTM (Hochreiter and Schmidhuber, 1997), Transformer (Vaswani et al., 2017), and Mamba (Gu and Dao, 2024). Table 10 presents a comparative analysis of their update dynamics, receptive fields, time complexity, and optimization properties.

As illustrated in Table 10, each backbone offers distinct trade-offs. RNN and LSTM offer linear complexity with respect to sequence length  $L$  but suffer from sequential computation bottlenecks and gradient decay. Transformers provide the strongest direct modeling of long-term dependencies via attention, yet their quadratic complexity  $\mathcal{O}(L^2)$  hinders scalability for extremely long histories. Mamba represents a compelling middle ground, achieving the linear scaling of RNN ( $\mathcal{O}(L)$ ) while approximating the expressive power of Transformers through input-dependent selection mechanisms. By supporting these diverse architectures, EST allows practitioners to tailor the balance between computational efficiency and modeling capacity based on specific deployment constraints.

## E End-to-End Workflow on ICEWS14

To make the formalization in Sections 3–4 more concrete, we present a worked example from ICEWS14 that instantiates Algorithm 1 on a real test query.

**Step 1. Query formulation.** Consider the test query:

(Refugee\_(Afghanistan), Make\_a\_visit, ?, 8016),

whose gold tail is Iran. EST treats this as a temporal forecasting query and retrieves the history of the head entity before timestamp  $t = 8016$ .

**Step 2. History retrieval and attention.** Table 11 lists the most attended events in the retrieved history window. The history shows a repeated interaction pattern with Iran under the relation Make\_a\_visit, and the most recent event receives the largest attention weight, consistent with temporal forecasting intuition.

**Step 3. Forward scoring.** Each retrieved event is encoded into a state-aware structural representation and then passed through the sequence backbone to produce contextualized hidden states. Query-specific attention aggregates them into a context vector  $\mathbf{c}_s$ , which is used to score candidate tails via  $\phi(\mathbf{c}_s, r, o)$ . In this example, repeated visits to Iran provide the dominant evidence for ranking the gold tail.

**Step 4. Persistent-state write-back.** After the forward pass, EST writes the resulting context back to the persistent states of the head entity through the fast and slow buffers. This closed-loop update allows future predictions involving Refugee\_(Afghanistan) to depend on accumulated temporal evidence rather than being reconstructed independently at each snapshot.

Table 10: Comparative analysis of sequential backbone instantiations within the EST framework. We contrast the core mathematical form, temporal mixing mechanism, theoretical receptive field, and computational complexity for each backbone.  $L$  denotes the sequence length, and  $d$  represents the hidden dimension.

Backbone	Core Mechanism	Temporal Mixing Operator $\mathcal{T}$	Receptive Field & Dependencies	Time Complexity (Train/Infer)	Gradient Flow & Stability
<b>RNN</b>	Non-linear Recurrence	$\mathbf{h}_t = \tanh(\mathbf{W}_u \mathbf{u}_t + \mathbf{W}_h \mathbf{h}_{t-1})$ (Implicit 1-step Markov mixing)	Infinite impulse response (theoretically); implies dependency on full history $\mathbf{u}_{<t}$ via recursion.	$\mathcal{O}(Ld^2)$ $\mathcal{O}(Ld^2)$	Susceptible to vanishing/exploding gradients over long $L$ ; lacks explicit gating stabilization.
<b>LSTM</b>	Gated Recurrence	$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t$ $\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t)$	Causal; the additive memory cell $\mathbf{c}_t$ creates a "constant error carousel" to retain long-term dependencies.	$\mathcal{O}(Ld^2)$ $\mathcal{O}(Ld^2)$	Gating mechanisms alleviate gradient decay; stable optimization for medium-to-long sequences.
<b>Transformer</b>	Global Attention	$\mathbf{A} = \text{softmax}\left(\frac{\mathbf{QK}^\top}{\sqrt{d}} + \mathbf{M}\right) \mathbf{V}$ (Content-based global kernel)	Global receptive field ( $L$ ); direct path between any pair $(t, \tau)$ within the window regardless of distance.	$\mathcal{O}(L^2d)$ $\mathcal{O}(L^2d)$	Shortest gradient path ( $\mathcal{O}(1)$ ); highly parallelizable but memory-intensive for large $L$ .
<b>Mamba</b>	Selective State Space	$\mathbf{s}_t = \bar{\mathbf{A}}_t \mathbf{s}_{t-1} + \bar{\mathbf{B}}_t \mathbf{u}_t$ $\mathbf{h}_t = \mathbf{C}_t \mathbf{s}_t$ (Input-dependent Linear Scan)	Causal; combines RNN-like recursion with CNN-like parallel scan. Selection mechanism compresses context adaptively.	$\mathcal{O}(Ld)$ $\mathcal{O}(Ld)$	Linear scalability with stable long-range propagation due to discretized state evolution.

Table 11: Worked example of history retrieval and query-specific attention on ICEWS14. For brevity, we show the most attended events in the retrieved history window.

Order	Historical event	$\Delta t$	Attention
1	(Refugee_(Afghanistan), Make_a_visit, Iran, 7944)	72	0.1753
2	(Refugee_(Afghanistan), Make_a_visit, Iran, 4728)	3288	0.1403
3	(Refugee_(Afghanistan), Make_a_visit, Iran, 5304)	2712	0.1412
4	(Refugee_(Afghanistan), Make_a_visit, Iran, 4848)	3168	0.1380
5	(Refugee_(Afghanistan), Make_a_visit, Iran, 4992)	3024	0.1344

**Takeaway.** This example illustrates how EST processes a real query end-to-end: query formulation, history retrieval, query-aware aggregation, tail scoring, and persistent state evolution.

---

**Algorithm 1:** Entity State Tuning (EST) for Temporal KG Extrapolation

---

**Input:** Training set  $\mathcal{D}$ ; entity/relation embeddings  $\mathbf{E}_{\mathcal{E}}, \mathbf{E}_{\mathcal{R}}$ ; global states ( $\mathbf{S}_{\text{fast}}, \mathbf{S}_{\text{slow}}$ ); structural encoder  $\Phi_{\text{struct}}$ ; sequence backbone  $\mathcal{T}$ ; time map  $\varphi(\cdot)$ ; hyperparams  $\lambda, \kappa, \gamma$ ; negative sampler  $\text{Neg}(\cdot)$ .

**Output:** Model parameters  $\Theta$  and updated global states ( $\mathbf{S}_{\text{fast}}, \mathbf{S}_{\text{slow}}$ ).

**for** each minibatch  $\mathcal{B} \subset \mathcal{D}$  **do**

$\mathcal{L} \leftarrow 0$

**for** each quadruple  $(s, r, o, t) \in \mathcal{B}$  **do**

        Construct history  $\mathcal{H}_s(t) = \{(o_k, r_k, \tau_k)\}_{k=1}^L$

        // (1) Topology-Aware State Perceiver

**for**  $k = 1$  to  $L$  **do**

$\mathbf{S}_{o_k} \leftarrow \mathbf{S}_{\text{slow}}(o_k)$

$\mathbf{g}^k \leftarrow \sigma(\mathbf{W}_g[\mathbf{E}_{o_k} \parallel \mathbf{S}_{o_k}])$

$\tilde{\mathbf{E}}_{o_k} \leftarrow \mathbf{g}^k \odot \mathbf{E}_{o_k} + (1 - \mathbf{g}^k) \odot \mathbf{S}_{o_k}$

$\mathbf{X}_k \leftarrow \Phi_{\text{struct}}(\tilde{\mathbf{E}}_{o_k}, \mathbf{E}_{r_k})$

$\mathbf{u}_k \leftarrow \text{MLP}([\mathbf{X}_k \parallel \mathbf{E}_{r_k} \parallel \varphi(t - \tau_k)])$

        // (2) Unified Temporal Context Module

$\mathbf{h}_{1:L} \leftarrow \mathcal{T}(\mathbf{u}_{1:L}), \quad \mathbf{y}_k \leftarrow \mathcal{P}(\mathbf{h}_k)$

        Compute query-aware attention  $\{\alpha_k\}_{k=1}^L$  and context  $\mathbf{c}_s \leftarrow \sum_{k=1}^L \alpha_k \mathbf{y}_k$

        // (3) De-confounded State Evolution

$\mathcal{E}^- \leftarrow \text{Neg}(s, r, t)$

$\mathcal{L} \leftarrow \mathcal{L} - \phi(\mathbf{c}_s, r, o) + \log \sum_{o' \in \{o\} \cup \mathcal{E}^-} \exp(\phi(\mathbf{c}_s, r, o'))$

        // Closed-loop state evolution (dual-system update)

$\mathbf{S}_{\text{fast}}(s) \leftarrow (1 - \lambda)\mathbf{S}_{\text{fast}}(s) + \lambda \mathbf{c}_s$

$\delta_s \leftarrow \|\mathbf{S}_{\text{fast}}(s) - \mathbf{S}_{\text{slow}}(s)\|_2$

$g_s \leftarrow \sigma(\kappa(\delta_s - \gamma))$

$\mathbf{S}_{\text{slow}}(s) \leftarrow \mathbf{S}_{\text{slow}}(s) + g_s(\mathbf{S}_{\text{fast}}(s) - \mathbf{S}_{\text{slow}}(s))$

    Update  $\Theta$  by minimizing  $\mathcal{L}$

---