

Negative-Aware Diffusion Process for Temporal Knowledge Graph Extrapolation

Yanglei Gan, Peng He, Yuxiang Cai, Run Lin, Guanyu Zhou, Qiao Liu*
University of Electronic Science and Technology of China
{yangleigan, hepenglk, yuxiangcai, runlin, 202522080824}@std.uestc.edu.cn,
qliu@uestc.edu.cn

Abstract

Temporal Knowledge Graph (TKG) reasoning seeks to predict future missing facts from historical evidence. While diffusion models (DM) have recently gained attention for their ability to capture complex predictive distributions, two gaps remain: (i) the generative path is conditioned only on positive evidence, overlooking informative negative context, and (ii) training objectives are dominated by cross-entropy ranking, which improves candidate ordering but provides little supervision over the calibration of the denoised embedding. To bridge this gap, we introduce **Negative-Aware Diffusion** model for TKG Extrapolation (NADEx). Specifically, NADEx encodes subject-centric histories of entities, relations and temporal intervals into sequential embeddings. NADEx perturbs the query object in the forward process and reconstructs it in reverse with a Transformer denoiser conditioned on the temporal-relational context. We further derive a cosine-alignment regularizer derived from batch-wise negative prototypes, which tightens the decision boundary against implausible candidates. Comprehensive experiments on four public TKG benchmarks demonstrate that NADEx delivers state-of-the-art performance¹.

1 Introduction

Temporal Knowledge Graphs (TKGs) serve as dynamic structures representing entities and their evolving relationships through time (Ji et al., 2021; Liang et al., 2024), expressed via quadruples of the form (s, r, o, t) , such as (*Donald Trump, meet, Joseph Biden, 2024/12/21*). Reasoning over TKGs predicts unknown links from historical observations and comprises two settings: **interpolation** (Cai et al., 2023; Luo et al., 2024), which infers missing facts within the observed timeline, and

extrapolation, which forecasts events beyond it (Yang et al., 2015; Zhu et al., 2021; Li et al., 2021). We focus on extrapolation as it provides anticipatory insight for future-oriented decision-making.

Accurate forecasting of future events necessitates a deep and nuanced comprehension of historical event sequences and relational dynamics. Prior research efforts (Jin et al., 2020; Li et al., 2021; Trivedi et al., 2017) commonly adopt deterministic embedding strategies by integrating neighboring structural contexts and temporal dependencies. These representations are subsequently utilized within scoring functions such as TransE (Bordes et al., 2013), DistMult (Yang et al., 2015), to assess the plausibility of prospective facts.

Recent progress in TKG reasoning has largely followed a **learn-to-classify** paradigm, seeking ever-richer representations while maintaining deterministic prediction targets. One prominent line of research augments graph neural networks (GNNs) with temporal gates that selectively transmit historical information, enabling fine-grained updates to entity and relation embeddings (Li et al., 2021, 2022b). A complementary stream adopts contrastive objectives, sharpening decision boundaries by contrasting local versus global neighborhoods (Chen et al., 2024b; Pang et al., 2025) or historical versus non-historical contexts (Xu et al., 2023). Rule-based techniques further enhance interpretability by injecting symbolic priors into embedding models (Chen et al., 2025b,c), while reinforcement-learning approaches guide multi-hop exploration with temporally calibrated rewards (Sun et al., 2021; Zheng et al., 2023).

Although these techniques secure strong empirical results, their deterministic nature limits their ability to model the inherent uncertainty of future events. Motivated by the generative prowess of diffusion models (DMs), recent work has reframed TKG reasoning as a **learn-to-generate** paradigm, exploiting DMs to sample plausible target object

*Corresponding Author

¹The source code is online at: <https://github.com/AONE-NLP/TKG-NADEx>.

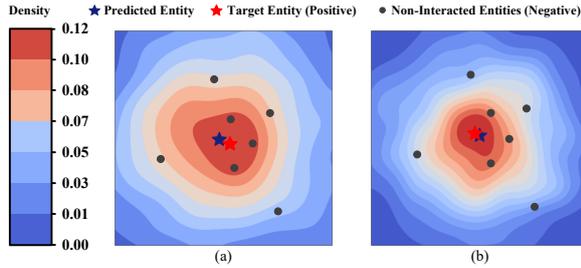


Figure 1: Illustration of entity embedding distributions learned by Diffusion-based TKG reasoning framework. (a) Neglecting the negative item distribution leads to predicted entities potentially drifting toward non-interacted negative entities. (b) Incorporating the negative sampling sharpens the predictive distribution, positioning the predicted entity centroid closer to the true target.

(Cai et al., 2024; Chen et al., 2025a; Cao et al., 2025). Despite their empirical success, existing methods exhibit two key limitations:

- **Neglect of Negative Context in Query Conditioning.** Existing DM-based methods generate entity embeddings by conditioning solely on observed positive facts, without accounting for negative examples that are essential for discrimination (Xu et al., 2023; Chen et al., 2024b; Cao et al., 2025). As illustrated in Figure 1(b), explicitly incorporating negative samples into the generative diffusion process introduces counterfactual evidence that regularizes the trajectory and reduces false-positive bias, yielding a lower-variance predictive distribution that is more tightly centered on true positives.
- **Reliance on Generic Loss Formulations.** Most existing work optimizes diffusion dynamics with generic reconstruction losses such as vanilla cross-entropy ranking (Cai et al., 2024), that emphasize relative ordering of candidate entities rather than enforcing explicit margins separating plausible from implausible outcomes (Chen et al., 2025a). Therefore, the generative potential of diffusion remains under-exploited.

In this paper, we propose a **Negative-Aware Diffusion** model for temporal knowledge graph **Extrapolation** (NADEx). Specifically, we reformulate temporal knowledge graph (TKG) reasoning as a sequence prediction task by encoding historical objects, event types, and temporal intervals into sequential embeddings. To capture uncertainty in future events, Gaussian noise is injected into target entity embeddings, and a Transformer-based denoising network conditioned on historical contexts and

temporal dynamics reconstructs these embeddings. Moreover, we develop a cosine-alignment regularizer for diffusion-based TKG reasoning that leverages batch-wise negative prototypes, computed by averaging the embeddings of the other targets in the mini-batch. Used alongside cross-entropy reconstruction loss, this direction-aware constraint sharpens separation between plausible and implausible candidate entities. Our contributions are three-fold:

- We propose NADEx, a negative-aware diffusion model for temporal knowledge graph extrapolation, explicitly capturing the dynamic and stochastic nature of future events through a Gaussian perturbed sequence denoising process.
- NADEx optimizes a cross-entropy reconstruction objective augmented with a cosine-alignment regularizer, sharpening separation between plausible and implausible candidate entities thus yielding a more discriminative predictive distribution.
- Extensive experiments conducted on four real world datasets demonstrate that NADEx yields superior performance, while remaining robust across diverse evaluation scenarios.

2 Related Works

2.1 Discriminative TKG Reasoning

Temporal Knowledge Graph (TKG) reasoning seeks to infer future facts by leveraging the chronological trajectory of previously observed triples. Early work such as Know-Evolve (Bordes et al., 2013; Chen et al., 2024a) treats the TKG as a set of mutually-exciting Hawkes processes, yielding continuous-time intensity functions for entity–relation interactions. THCN (Chen et al., 2024a) couples a Hawkes process with a temporal causal convolutional network, enabling causal convolutions to inherit the self-exciting inductive bias of point processes. However, these methods propagate information only to adjacent timestamps, leaving long-range dependencies under-exploited. Subsequent studies enhance static KG encoders with temporal signals. RE-NET (Jin et al., 2020), RE-GCN (Li et al., 2021), and CyGNet (Zhu et al., 2021) inject recurrent gating, temporal self-attention, and cycle-aware constraints, respectively, to capture evolving patterns, while structural extensions, CEN (Li et al., 2022c), xERTE (Han et al., 2021) and HisMatch (Li et al., 2022d) exploit hierarchical alignment for richer representations. Yet

their reliance on single-scale memories still limits sensitivity to distant context.

Considering the long-term dependencies among entities and relations, Zhang et al. (Zhang et al., 2023) decompose each entity–relation trajectory into long- and short-term streams. Building on that idea, CENET (Xu et al., 2023) formulates temporal reasoning as a history-aware contrastive task: for every query, positive events are contrasted against hard negatives drawn from both historical and non-historical contexts. LogCL (Chen et al., 2024b) further refines this idea with a local–global paradigm that first selects query-relevant snapshot windows via entity-centric attention and then applies four bespoke contrastive patterns.

Parallel to these efforts, symbolic and structural approaches pursue interpretability and inductive generalization. Rule-based systems such as Tlogic (Liu et al., 2022) mine time-stamped first-order clauses via temporal random walks, offering interpretable but coverage-limited priors. DaeMon (Dong et al., 2023) aggregate relation-specific temporal paths between query entities, replacing brittle entity embeddings with structural evidence that generalizes to unseen nodes. CognTKE (Chen et al., 2025c) further unifies shallow reasoning and local deep path reasoning over a cognitive temporal relation digraph, achieving notable gains.

2.2 Generative TKG Reasoning

Most recently, generative frameworks have introduced principled uncertainty modelling. DiffuTKG (Cai et al., 2024) recasts forecasting as conditional sequence denoising under a forward-and-reverse Gaussian diffusion process, and DPCL-Diff (Cao et al., 2025) couples graph-node diffusion with dual-domain periodic contrastive learning to disentangle recurrent cycles from genuinely novel events. Complementing these methods, Luo et al. (Luo et al., 2024) leverage large language models (LLMs) to generate multi-step event chains that capture higher-order dependencies, while LLM-DR (Chen et al., 2025a) harnesses a classifier-free guided diffusion model conditioned on temporal KG contexts and refines generated candidate rules with LLM-based constraints. Despite these advances, existing methods often overlook negative contexts and rely on generic loss formulations, resulting in blurred predictive distributions. Such limitations underscore the need for a generative framework that can sharply discriminate future events.

3 Preliminary

Definition 1. Temporal Knowledge Graph. Let \mathcal{E} , \mathcal{R} , and \mathcal{T} denote finite sets of entities, relation types, and timestamps, respectively. A temporal knowledge graph \mathcal{G} is a collection of time-stamped quadruples:

$$\mathcal{G} = \left\{ (s, r, o, t) \mid s, o \in \mathcal{E}, r \in \mathcal{R}, t \in \mathcal{T} \right\}, \quad (1)$$

where each tuple encodes the fact that relation r holds from subject s to object o at time t . More specifically, the TKG can be viewed as an ordered sequence of static snapshots:

$$\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_{|\mathcal{T}|}\}, \quad \mathcal{G}_t \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{T}, \quad (2)$$

where \mathcal{G}_t aggregates all triples that are valid at timestamp t . Following the standard bidirectional–relation convention (Kazemi and Poole, 2018), we augment every quadruple (s, r, o, t) with its inverse (o, r^{-1}, s, t) , where r^{-1} is a distinct relation denoting the reverse semantics of r .

Definition 2. Temporal Knowledge Graph Reasoning. Let $q = (s, r, ?, t)$ be a query quadruple whose object entity is missing at timestamp t . Given the sliding history window of length L , $\mathcal{G}_{t-L-1:t-1} = \{\mathcal{G}_{t-L}, \mathcal{G}_{t-L+1}, \dots, \mathcal{G}_{t-1}\}$. The TKG reasoning task is to learn a scoring function:

$$\text{score}_t(o) = f(s, r, o, \mathcal{G}), \hat{o} = \arg \max_{o \in \mathcal{E}} \text{score}_t(o). \quad (3)$$

Each candidate object $o \in \mathcal{E}$ is assigned a plausibility score and the highest-scoring entity completes the quadruple.

4 Method

4.1 Temporal Representation Learning

Given a query $q = (s, r, ?, t)$, we follow Cai et al. (2024) and cast TKG extrapolation as sequence prediction over the subject’s history. Let the subject-centric history up to time $t - 1$ be $Q_{0:t-1} = \{(s, r_0, o_0, t_0), \dots, (s, r_i, o_i, t_i), \dots, (s, r_{t-1}, o_{t-1}, t_{t-1})\}$, where L is the window length. We form three aligned sequences $O_{0:t-1}$, $R_{0:t-1}$, $T_{0:t-1}$, and employ three trainable matrices to get their corresponding embeddings,

$$\begin{aligned} \mathbf{O}^+ &= \mathbf{E}_o [o_0; o_1; \dots; o_{t-1}, o_t]; \\ \mathbf{R} &= \mathbf{E}_r [r_0; r_1; \dots; r_{t-1}, r_t]; \\ \mathbf{T} &= \mathbf{E}_{\Delta t} [t_0; t_1; \dots; t_{t-1}, t_t]. \end{aligned} \quad (4)$$

where $\mathbf{E}_o \in \mathbb{R}^{\mathcal{E} \times h}$, $\mathbf{E}_r \in \mathbb{R}^{\mathcal{R} \times h}$, $\mathbf{E}_{\Delta t} \in \mathbb{R}^{t \times h}$. \mathcal{E} , \mathcal{R} , and h denotes the number of entities, event

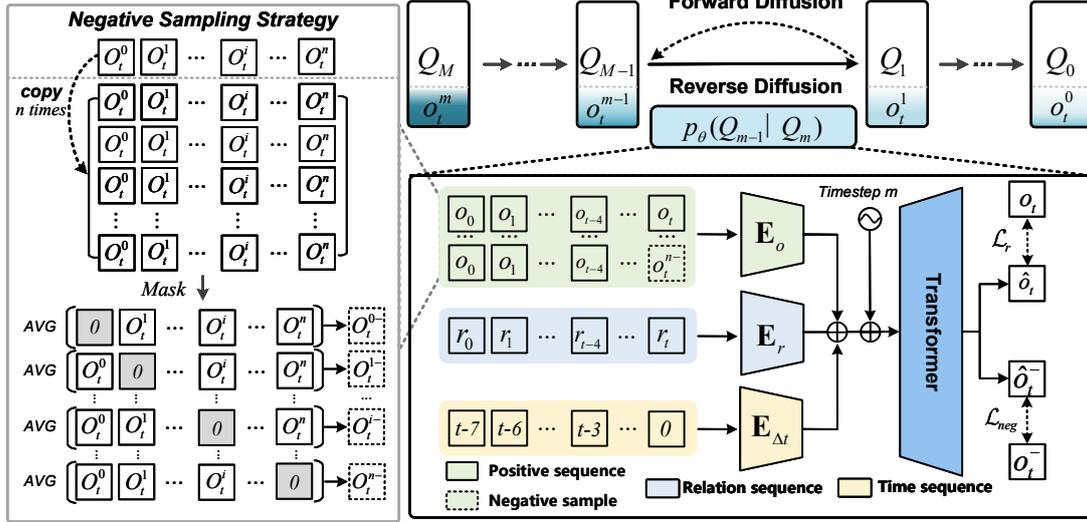


Figure 2: Overview of NADEX. In Forward Diffusion, Gaussian noise is applied to the future object o . During the reverse denoising process, NADEX incorporates both positive and negative (o_t^-) samples to disentangle plausible from implausible event predictions. For negative samples, we first stack all target embeddings in the mini-batch into a single row, copy this row n times to form an $n \times n$ matrix, set the diagonal to 0 to exclude each target itself, and then take the row-wise mean of the remaining entries to obtain one compact negative prototype per target.

types and the embedding dimension, respectively. $\mathbf{O}_{0:t-1}$ denotes the embeddings of the historical object entities, denoted as $\mathcal{H} = \mathbf{O}_{0:t-1}$ for brevity, while $\mathbf{o}_t \in \mathbb{R}^{1 \times h}$ refers to the representation of the target entity to be predicted.

4.2 Negative Sampling Strategy

To exploit the discriminative signal from negative evidence, we design a batch-wise negative sampling strategy. For a mini-batch of size n , we first collect the target entities into a sequence $\mathbf{O}_t = \{o_t^0, o_t^1, \dots, o_t^n\} \in \mathbb{R}^{n \times 1}$, and then construct an $n \times n$ mask by permuting the diagonal entries, resulting in each row containing the other $N - 1$ objects from the batch as negatives. **Averaging** these negative candidates along each row generates a compact negative sequence, defined as:

$$\begin{pmatrix} o_t^0 & o_t^1 & \dots & o_t^n \\ o_t^0 & o_t^1 & \dots & o_t^n \\ \vdots & \vdots & \ddots & \vdots \\ o_t^0 & o_t^1 & \dots & o_t^n \end{pmatrix} \Rightarrow \text{AVG.} \begin{pmatrix} (0 & o_t^1 & \dots & o_t^n) \\ (o_t^0 & 0 & \dots & o_t^n) \\ (\vdots & \vdots & \ddots & \vdots) \\ (o_t^0 & o_t^1 & \dots & 0) \end{pmatrix} = \begin{pmatrix} o_t^{0-} \\ o_t^{1-} \\ \vdots \\ o_t^{n-} \end{pmatrix}. \quad (5)$$

The treatment of negative samples o_t^- is equivalent to o_t , which is obtained through the entity embedding matrix \mathbf{E}_o :

$$\mathbf{o}_t^- = \mathbf{E}_o[o_t^-]. \quad (6)$$

Finally, our negative sequence can be expressed as:

$$\mathbf{O}^- = [\mathbf{o}_0; \mathbf{o}_1; \dots; \mathbf{o}_{t-1}, \mathbf{o}_t^-]. \quad (7)$$

It is important to note that, due to the event-stream formulation of temporal knowledge graph (TKG) extrapolation and the way events are organized in our benchmarks, we treat all events sharing the same timestamp as a single training unit. Concretely, if N events occur at a given timestamp, we construct one full batch that contains all N corresponding target entities. This design follows directly from the standard TKG reasoning objective: at timestamp $t + 1$, the model is required to predict the missing object entities for all queries occurring at that time. Consequently, N is not a tunable batch-size hyperparameter but is uniquely determined by the number of real-world events recorded at each timestamp, which can vary substantially across time.

4.3 Forward Diffusion Process

We model predictive uncertainty by diffusing only the candidate object embedding \mathbf{o}_t and sampled negatives \mathbf{o}_t^- . The forward process progressively add Gaussian noise over M steps:

$$\begin{aligned} o_m^+ &= \sqrt{\bar{\alpha}_m} o_t + \sqrt{1 - \bar{\alpha}_m} \epsilon, \\ o_m^- &= \sqrt{\bar{\alpha}_m} o_t^- + \sqrt{1 - \bar{\alpha}_m} \epsilon, \end{aligned} \quad (8)$$

where $\epsilon \in \mathcal{N}(0, \mathbf{I})$ introduces Gaussian noise, and $\bar{\alpha}_m \in (0, 1)$ controls the signal-to-noise ratio at step m , simulating the predictive uncertainty.

To regulate the accumulation of noise over time, we employ a linear schedule that defines $1 - \bar{\alpha}_m$ as:

$$1 - \bar{\alpha}_m = \delta \cdot \left(\alpha_{\min} + \frac{m-1}{M-1} (\alpha_{\max} - \alpha_{\min}) \right). \quad (9)$$

where $\delta \in [0, 1]$ is a global scaling factor that moderates the overall diffusion strength, and α_{\min} and α_{\max} are hyperparameters that determine the lower and upper bound of noise levels.

4.4 Reverse Denoising Process

In reverse process, we aim to recover the target sequence of object entities representation o_t or o_t^- iteratively from a pure Gaussian noise o_m^+ or o_m^- . At each step, the model is conditioned on the temporal **T** and relational **R** context:

$$p_{\theta}(\hat{o}_{m-1} | *) = \mathcal{N}(\hat{o}_{m-1}; \mu_{\theta}(*), \Sigma_{\theta}(*)). \quad (10)$$

where we use a Transformer to model $\mu_{\theta}(*)$ and $\Sigma_{\theta}(*)$ during reverse process. For brevity, $*$ denotes $[\hat{o}_m, \mathbf{r}, \mathbf{t}, m]$, which is the denoising conditions. \hat{o}_m is set to \mathcal{H} for the first step:

$$\begin{aligned} \hat{o}_0^+ &= f_{\theta}(\mathbf{Z}_x) = \text{Transformer}([\mathbf{o}_i^+ + \text{Emb}(i)]), \\ \hat{o}_0^- &= f_{\theta}(\mathbf{Z}_x) = \text{Transformer}([\mathbf{o}_i^- + \text{Emb}(i)]). \end{aligned} \quad (11)$$

where $i \in [0, m]$, $\text{Emb}(\cdot)$ denotes denoising step embeddings to manage the hidden representations at different noise levels.

4.5 Training Objective

Reconstruction Loss. We introduce a distance quantification measurement based on the dot product between embeddings:

$$\mathbf{Y} = \text{Softmax}(\hat{\mathbf{o}}_0^+ \cdot (\mathbf{E})^T), \quad (12)$$

where $\hat{\mathbf{o}}_0^+ \in \mathbb{R}^{1 \times h}$ denotes the representation of the target object from the reverse process. $(\cdot)^T$ is the matrix transposition operation. Moreover, we utilize a reconstruction loss function \mathcal{L}_r as follows:

$$\mathcal{L}_r = - \sum_{i \in \{1, 2, \dots, d\}} y_i \log(\mathbf{Y}_i), \quad (13)$$

where y_i denotes the one-hot encoding of the i -th gold entity, and \mathbf{Y}_i is the predicted probability.

Negative Sampling Loss. Exclusively depending on the generation of optimization objectives may result in model overfitting to historical frequent events. Though this method can alleviate the impact of sparse and noisy interactions inherent in the system, it might fail to fully harness the generative

and generalization capacities of diffusion models, thereby leading to inaccurate assessments of both unseen and observed facts. To gain a better understanding and to specifically redesign for modeling object distributions in order to achieve diffusion optimization goals. Specifically, our negative sample cosine loss (\mathcal{L}_{neg}) is represented as:

$$\mathcal{L}_{\text{neg}} = \frac{1}{N} \sum_{i=1}^N \left(\frac{\mathbf{o}_i^-}{\|\mathbf{o}_i^-\|_2} \cdot \frac{\hat{\mathbf{o}}_i^-}{\|\hat{\mathbf{o}}_i^-\|_2} - 1 \right)^2. \quad (14)$$

The training objective can be formulated as:

$$\mathcal{L} = -(1 - \lambda) \log \sigma(-\gamma \cdot (\mathcal{L}_r - \mathcal{L}_{\text{neg}}) + \epsilon) + \lambda \mathcal{L}_r. \quad (15)$$

where σ is the Sigmoid function, $\epsilon = 1e - 8$ to prevent numerical instability, and λ and γ are equilibrium weight hyperparameters.

4.6 Inference

During the inference phase, the denoising network $f_{\theta}(*)$ is employed to reconstruct the original samples from Gaussian noise ϵ_m through a progressive refinement process. The denoising generation process of target \hat{o}_0 is formulated as:

$$\begin{aligned} \hat{\mathbf{o}}_m &= [\mathbf{o}_{0:m}; \epsilon_m] + \mathbf{r}_t + \mathbf{t}_t, \\ \hat{o}_0 &= f_{\theta}(\hat{\mathbf{o}}_m, m). \end{aligned} \quad (16)$$

Subsequently, we compute the similarity between \hat{o}_0 and the object embedding matrix $\mathbf{E}_o \in \mathbb{R}^{|\mathcal{E}| \times h}$ to derive ranking results:

$$P = \hat{o}_0 \cdot \mathbf{E}_o^T. \quad (17)$$

Entities are then ranked in descending order based on these similarity scores P . The top-K ranked entities are evaluated using standard metrics such as Mean Reciprocal Rank (MRR) and Hit@K.

5 Experiments

5.1 Experimental Setups

Datasets. Our experiments employ four benchmark datasets, including ICEWS14, ICEWS05-15, ICEWS18, and GDEL, to evaluate the proposed model. Specifically, the ICEWS datasets originate from the Integrated Crisis Early Warning System (Boschee et al., 2015), while the GDEL dataset is sourced from the Global Database of Events, Language, and Tone (Leetaru and Schrodt, 2013). The data statistics are summarized in Appendix C.1.

Baseline Models. We benchmark NADEX against a broad set of state-of-the-art knowledge-graph reasoning methods across three categories: **Static**

Table 1: Overall performance comparison (%) on four datasets in terms of MRR and Hit@1/3/10 with time-aware metrics. The best results are highlighted in **bold**, while the second-best results are underlined. Statistical significance is determined via paired bootstrap t -test, with differences at $p < 0.05$ deemed significant.

Models	ICEWS14				ICEWS18				ICEWS05-15				GDEL T			
	MRR	Hit@1	Hit@3	Hit@10												
DisMult (2015)	15.44	10.91	17.24	23.92	11.51	7.03	12.87	20.86	17.95	13.12	20.71	29.32	8.68	5.58	9.96	17.13
ConvE (2018)	35.09	25.23	39.38	54.68	24.51	16.23	29.25	44.51	33.81	24.78	39.00	54.95	16.55	11.02	18.88	31.60
RotatE (2019)	21.31	10.26	24.35	44.75	12.78	4.01	14.89	31.91	24.71	13.22	29.04	48.16	13.45	6.95	14.09	25.99
TTransE (2018)	13.72	2.98	17.70	35.74	8.31	1.92	8.56	21.89	15.57	4.80	19.24	38.29	5.50	0.47	4.94	15.25
TA-DistMult (2018)	25.80	16.94	29.74	42.99	16.75	8.61	18.41	33.59	24.31	14.58	27.92	44.21	12.00	5.76	12.94	23.54
DE-SimIE (2020)	33.36	24.85	37.15	48.92	19.30	11.53	21.86	34.80	35.02	25.91	38.99	52.75	19.70	12.22	21.39	33.70
RE-NET (2020)	36.93	26.83	39.51	54.78	28.81	19.05	32.44	47.51	43.32	33.43	47.77	63.06	19.62	12.42	21.00	34.01
RE-GCN (2021)	40.39	30.66	44.96	59.21	30.58	21.01	34.34	48.75	48.03	37.33	53.85	68.27	19.64	12.42	20.90	33.69
CyGNet (2021)	35.05	25.73	39.01	53.55	24.93	15.90	28.28	42.61	36.81	26.61	41.63	56.22	18.48	11.52	19.57	31.98
TITer (2021)	41.73	32.74	46.46	58.44	29.98	22.05	33.46	44.83	47.69	37.95	52.92	65.81	15.46	10.98	15.61	24.31
CEN (2022c)	42.20	32.08	47.46	61.31	31.50	21.70	35.44	50.59	46.84	36.38	52.45	67.01	20.39	12.96	21.77	34.97
TIRGN (2022b)	44.04	33.83	48.95	63.84	33.66	23.19	37.99	54.22	50.04	39.25	56.13	70.71	21.67	13.63	23.27	37.60
HisMatch (2022d)	46.42	35.91	51.63	<u>66.84</u>	33.99	23.91	37.90	53.94	<u>52.85</u>	42.01	59.05	73.28	22.01	14.45	23.80	36.61
RETIA (2023b)	42.76	32.28	47.77	<u>62.75</u>	32.43	22.23	36.48	52.94	47.26	36.64	52.90	67.76	20.12	12.76	21.45	34.49
CENET (2023)	39.02	29.62	43.23	57.49	27.85	18.15	31.63	46.98	41.95	32.17	46.93	60.43	20.23	12.69	21.70	34.92
CRAFT (2024)	45.71	35.05	51.83	65.21	34.21	23.96	38.53	54.11	50.14	39.56	56.18	70.09	23.78	<u>15.38</u>	<u>26.23</u>	<u>40.15</u>
THCN (2024a)	45.39	<u>36.58</u>	50.84	66.07	35.63	24.90	39.26	56.76	51.94	40.32	57.79	72.18	23.46	15.18	25.21	39.03
DiffuTKG (2024)	<u>47.58</u>	36.38	<u>53.41</u>	66.01	<u>35.65</u>	<u>27.19</u>	<u>44.04</u>	59.55	48.97	39.80	56.92	69.84	21.87	14.43	23.68	36.05
LogiQ (2025b)	44.71	35.72	51.03	64.21	34.94	24.76	39.57	56.32	51.04	40.71	57.55	71.00	-	-	-	-
CognTKE (2025c)	46.06	36.49	51.11	64.49	35.24	25.21	39.93	54.71	53.13	<u>42.62</u>	<u>59.42</u>	<u>72.70</u>	-	-	-	-
NADEx	49.03	39.45	57.48	70.55	36.84	27.58	45.12	60.58	52.17	43.38	60.75	71.93	<u>23.67</u>	16.10	28.35	43.05

methods: DisMult (Yang et al., 2015), ConvE (Dettmers et al., 2018), RotatE (Sun et al., 2019); **Interpolation methods:** TTransE (Leblay and Chekol, 2018), TA-DistMult (Garcia-Duran et al., 2018), DE-Simple (Goel et al., 2020); **Extrapolation methods:** RE-NET (Jin et al., 2020), Re-GCN (Li et al., 2021), TITer (Sun et al., 2021), CEN (Li et al., 2022c), TIRGN (Li et al., 2022b), HisMatch (Li et al., 2022d), RETIA (Liu et al., 2023b), CENET (Xu et al., 2023), CRAFT (Zhang et al., 2024), THCN (Chen et al., 2024a), DiffuTKG (Cai et al., 2024), LogiQ (Chen et al., 2025b), CognTKE (Chen et al., 2025c). We provide baseline descriptions in Appendix C.2.

Evaluation Metrics. To measure temporal extrapolation performance, we cast the task as masked entity prediction, where either the subject or object is held out in quadruples of the form $(s, r, ?, t)$ or $(?, r, o, t)$. Predictions are scored and ranked, and we report Mean Reciprocal Rank (MRR) alongside Hits@1, Hits@3, and Hits@10. All results are computed under the time-aware filtering protocol.

Implementation Details. We train all models with the Adam optimizer. For ICEWS14 and ICEWS18 we use a learning rate of $1e-3$ and $1e-5$ for ICEWS05-15 and GDEL T. Each model is trained for 100 epochs with an embedding dimension of 200 for both entities and relations and a dropout rate of 0.2 applied throughout. All experiments are executed on a single NVIDIA A100 (80G) GPU. All evaluations follow the filtered ranking protocol.

Statistical significance is assessed using a bootstrap paired t -test. In practice, our framework can be effectively trained within 12 hours on a single GPU on the largest dataset GDEL T.

5.2 Overall Performance

Table 1 summarizes NADEx’s performance against state-of-the-art (SOTA) baselines across four benchmark datasets. From these results, we make the following key observations:

- **Extrapolation-oriented methods uniformly outperform both interpolation and time-agnostic (“static”) approaches.** Static models, by design, ignore any temporal ordering and therefore cannot track how entity–relation interactions evolve over time. Interpolation models improve upon static baselines by filling in gaps within the observed timeline, but their training objective and architecture remain tethered to historical fact completion rather than forward-looking prediction. In contrast, extrapolation models explicitly learn the dynamics of temporal progression, enabling them to predict future facts.
- **Across all four benchmarks, NADEx consistently surpasses SOTA methods on nearly every metric.** In terms of MRR, NADEx improves over the runner-up by 3.05% on ICEWS14 and 3.33% on ICEWS18 and falls slightly short on ICEWS05-15 and GDEL T. Moreover, NADEx

Table 2: Ablation study results ICEWS14 and GDELT datasets in terms of MRR and Hit@1/10.

Settings	ICEWS14			GDELT		
	MRR	Hit@1	Hit@10	MRR	Hit@1	Hit@10
NADEx	49.03	39.45	70.55	23.67	16.10	43.05
w/o. E_{Δ}	43.16	33.35	62.06	18.82	12.26	36.32
w/o. E_r	40.05	30.23	58.47	18.21	11.76	35.52
w/o. $E_r \& E_{\Delta}$	32.48	22.71	51.69	18.06	11.66	35.21
w/o. neg	47.42	36.39	67.61	21.25	14.22	42.34

delivers substantial uplifts in Hit@K on GDELT dataset, recording performance gain of 4.68% at Hit@1, 8.08% at Hit@3, and 7.22% at Hit@10.

- **NADEx achieves superior results on all four benchmarks compared against the leading diffusion-based method DiffuTKG.** Notably, the performance gap widens on the more complex ICEWS05-15 and GDELT datasets. NADEx improves MRR by 6.53% on ICEWS05-15 and 8.23% on GDELT. We attribute this pattern to the richer, longer-range temporal interactions present in ICEWS05-15 and GDELT, which pose greater challenges for diffusion pipelines. By incorporating negative-aware conditioning and a likelihood-driven training objective, NADEx more effectively models these intricate dynamics, whereas DiffuTKG’s reconstruction loss struggles to discriminate plausible from spurious future facts under such complex scenarios.

5.3 Ablation Studies

To quantify the contribution of each NADEx component, we perform ablations on ICEWS14 and GDELT, measuring MRR, Hit@1, and Hit@10, as presented in Table 2. Removing the time-interval embedding (“w/o E_{Δ} ”) degrades MRR by 5.87 points on ICEWS14 and 4.85 points on GDELT, highlighting the importance of explicit temporal cues. Omitting the relation-type embedding (“w/o E_r ”) incurs an even larger drop across all metrics, confirming that relational context is essential for accurate reconstruction. Omitting both temporal and relational embeddings simultaneously (“w/o $E_r \& E_{\Delta}$ ”) Omitting both temporal and relational embeddings simultaneously. Lastly, eliminating our negative-sampling strategy (“w/o neg”) consistently lowers performance, affirming the value of negative-aware training in improving the model’s discriminative capability between plausible and

Table 3: Performance of predicting unseen events in terms of MRR and Hit@1 on ICEWS14 and ICEWS18.

Models	ICEWS14		ICEWS18	
	MRR	Hit@1	MRR	Hit@1
RE-GCN	23.26	13.91	15.08	7.09
CEN	22.06	13.28	15.41	8.20
RETIA	24.17	14.67	16.62	9.08
HisMatch	27.49	19.04	17.51	11.13
DiffuTKG	25.22	15.23	16.48	8.84
LogCL	<u>29.19</u>	18.72	<u>18.40</u>	<u>11.74</u>
NADEx	30.85	20.08	19.52	12.17
<i>Improve.</i>	5.89%	5.46%	6.09%	3.66%

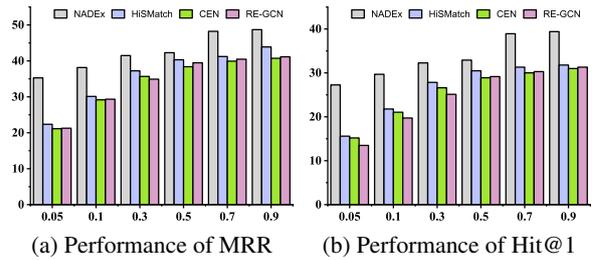


Figure 3: Performance of NADEx compared with three baselines in terms of MRR and Hit@1 on ICEWS14 under different training data scale settings.

implausible event predictions. Together, these results confirm that NADEx’s strength derives not merely from its generative backbone but from the harmonious integration of temporal, relational, and negative-aware conditioning.

5.4 Effectiveness of Unseen Events

To further verify the capacity of NADEx in capturing discriminative and generalizable representations for unseen event queries, we specifically evaluate its predictive performance on ICEWS14 and ICEWS18 datasets, both characterized by a relatively higher proportion of previously unobserved events. As reported in Table 3, NADEx consistently surpasses existing state-of-the-art methods by significant margins. Specifically, NADEx achieves remarkable improvements of approximately 4.43% and 5.40% in terms of absolute MRR on ICEWS14 and ICEWS18, respectively, compared with the strongest baseline (LogCL). Additionally, NADEx attains notable enhancements of 3.10% and 3.43% in absolute Hit@1 performance on the two datasets, underscoring its superior precision in accurately identifying true future entities. These empirical results substantiate that explicitly incorporating negative contexts and employing a diffusion-based generative framework effectively enhances the model’s discriminative power, enabling robust extrapolation

Table 4: Case studies on ICEWS14 showing each model’s top-5 predicted entities for sample queries.

Query	NADEX	LogCL	DiffuTKG	THCN
<China, 057, ?> Date: 2014-12-3 Label: Maldives	Maldives, 0.534	Japan, 0.457	Laos, 0.324	Laos, 0.298
	Taavi Roivas, 0.167	South Korea, 0.195	Thailand, 0.212	Malaysia, 0.141
	Angola, 0.019	Thailand, 0.081	Japan, 0.073	Iran, 0.102
	South Korea, 0.010	Legislature, 0.043	Iran, 0.053	Thailand, 0.095
	Malaysia, 0.006	Malaysia 0.035	Malaysia, 0.043	Japan, 0.074
<South Korea, 060, ?> Date: 2014-12-9 Label: Japan	Japan, 0.414	Japan, 0.289	North Korea (Military), 0.128	North Korea (Military), 0.329
	North Korea, 0.347	North Korea, 0.128	China, 0.122	North Korea, 0.122
	China, 0.084	North Korea, 0.027	Qatar, 0.078	China, 0.113
	Philippines, 0.028	Japan (Newspaper), 0.021	Iran, 0.060	Iran, 0.061
	Iran, 0.026	China, 0.006	Japan, 0.025	South Korea (Police), 0.025

to rare/unseen future events under uncertainty.

5.5 Performance under Limited Data

To assess data-efficiency, we subsample the ICEWS14 training set at several fractions (5%, 10%, 30%, 50%, 70%, and 90%) while keeping the validation and test splits intact. Figure 3 shows that every model deteriorates as the available data diminish, reflecting the increased difficulty of learning reliable representations under sparse supervision. Even so, NADEX maintains a clear lead over HisMatch, CEN, and RE-GCN in both MRR and Hit@1 at every sampling ratio. Notably, between the 50% and 5% splits, NADEX’s MRR degrades by only 7%, whereas all baselines suffer losses exceeding 15%. A similar trend holds for Hit@1. We attribute this robustness to the negative-aware diffusion mechanism, which supplies informative counter-examples during generation, the model learns fine-grained decision boundaries with far fewer labeled instances, and it maintains calibrated uncertainty estimates that remain reliable even when supervision is extremely sparse. Notably, the performance margin in favor of NADEX widens as the training fraction decreases, underscoring its superior generalization in low-resource settings. We attribute these strengths to NADEX’s negative-aware diffusion mechanism: by injecting negatives drawn from each mini-batch into the generation process, the model learns more discriminative decision boundaries with fewer samples.

5.6 Case Study

Table 4 presents two illustrative queries on ICEWS14, comparing NADEX against LogCL and DiffuTKG by listing each model’s top-5 predicted entities and their associated scores. For the first query <China, 057, ?, 2014-12-3>, NADEX places the correct target, **Maldives** (score 0.534), firmly

Table 5: Computational efficiency test on ICEWS14/18.

Model	ICEWS14		ICEWS18	
	Inf. Time	Params.	Inf. Time	Params.
RE-GCN	28.96s	26.52Mb	393.67s	42.68Mb
TiRGN	32.88s	43.35Mb	159.90s	59.59Mb
DiffuTKG	13.49s	18.35Mb	135.80s	34.46Mb
NADEX	10.91s	16.30Mb	96.95s	32.42Mb

at the top, followed by semantically coherent but less plausible alternatives such as Taavi Roivas (0.167) and Angola (0.019). LogCL instead places Japan first (0.457) and omits Maldives from its top-5; DiffuTKG similarly favors Laos (0.324) and does not rank Maldives among its candidates. For the second query <South Korea, 060, ?, 2014-12-9>, NADEX correctly identifies **Japan** as the most likely future partner (0.414), with close runner-up North Korea (0.347) reflecting plausible regional dynamics. LogCL also selects Japan but with lower confidence (0.289) and a more diffuse allocation across other entities; DiffuTKG ranks North Korea highest (0.128) and places Japan fifth (0.025). Across both queries, NADEX concentrates probability mass on the true target and suppresses near-neighbor distractors. The negative-aware conditioning sharpens the posterior and yields more calibrated confidence relative to the baselines.

5.7 Computational Efficiency

On computational efficiency, NADEX attains the best speed–size trade-off on both ICEWS14 and ICEWS18, as shown in Figure 5. It yields the lowest inference time and smallest parameter 10.91s/16.30Mb on ICEWS14 and 96.95s/32.42Mb on ICEWS18. Relative to the diffusion baseline DiffuTKG, NADEX reduces inference time by

Table 6: Performance comparison (%) with LLM-driven approaches on ICEWS14 and ICEWS18. The best results are highlighted in **bold**, while the second-best results are underlined. The experimental results are retrieved from (Luo et al., 2024) and (Chen et al., 2025a), respectively.

Models		GenTKG (2024)	GPT-NeoX (2022)	Llama-ICL (2023)	Vicuna-ICL (2023)	Llama-CoH (2024)	Vicuna-CoH (2024)	LLM-DR (2025a)	NADEx
ICEWS14	Hit@1	34.90	29.50	28.60	28.10	34.90	32.80	40.60	39.45
	Hit@3	47.30	40.60	39.70	39.10	47.00	45.70	<u>55.80</u>	57.48
	Hit@10	61.90	47.50	45.30	45.30	59.10	65.60	<u>67.00</u>	70.55
ICEWS18	Hit@1	21.50	17.70	17.70	17.20	22.30	20.90	30.40	27.58
	Hit@3	36.60	29.00	29.40	28.80	36.30	34.70	40.90	45.12
	Hit@10	49.60	38.20	36.40	36.40	52.20	53.60	<u>55.60</u>	60.58
AVG.		41.97	35.42	32.85	32.48	41.97	42.22	<u>48.38</u>	50.13

19.1% on ICEWS14 and 28.6% on ICEWS18 with smaller parameter sizes. Against graph-based baselines, NADEx is much faster than RE-GCN/TiRGN on ICEWS14 ICEWS18, while using 20% to 60% fewer parameters. When scaling from ICEWS14 to ICEWS18, the inference time of NADEx grows by 8.89 \times , compared with 10.07 \times for DiffuTKG and 13.59 \times for RE-GCN, indicating more favorable scaling with dataset size. These results demonstrate that NADEx achieves better computational efficiency without increasing model capacity.

5.8 Comparison with LLM-based Approaches

In this section, we systematically evaluate our proposed approach against recent state-of-the-art large language model (LLM)-driven temporal knowledge graph (TKG) reasoning methods. Specifically, we benchmark our model against prominent LLM-based baselines including GenTKG (Liao et al., 2024), GPT-NeoX (Lee et al., 2023), Llama-2 (Luo et al., 2024), Vicuna (Luo et al., 2024), and LLM-DR (Wang et al., 2024) on two benchmarks, ICEWS14 and ICEWS18. Performance is measured across standard ranking metrics: Hit@1, Hit@3, and Hit@10, and Table 6 summarizes the comparative results. Experimental results indicate that certain LLM-based approaches such as GenTKG and GPT-NeoX do not consistently outperform traditional TKG reasoning methods. This observation can be primarily attributed to the overly general logical rules derived from LLM-generated reasoning chains, which may fail to precisely model the intricate temporal dynamics inherent in TKGs. Notably, LLM-DR, which continuously refines its rule set to mirror the empirical distribution, yields the best Hit@1 scores (40.60% on ICEWS14 and 30.40% on ICEWS18 datasets). However, its Hit@3 and Hit@10 scores decline

because the iterative rule-pruning aggressively narrows each query to a high-precision candidate set, increasing the likelihood that the top prediction is correct while simultaneously sacrificing lower-rank diversity and overall recall. In contrast, our proposed approach achieves robust performance across all metrics, surpassing all baseline methods on Hit@3 and Hit@10, and obtaining a strong Hit@1 result second only to LLM-DR. With an average overall hit score of 50.13% across both datasets.

6 Conclusion

In this paper, we propose NADEx, a novel Negative-Aware Diffusion model specifically designed for TKG extrapolation. By recasting TKG reasoning as a sequence-denoising problem, NADEx perturbs target-entity embeddings with Gaussian noise and then leverages a Transformer-based denoiser to reconstruct the true entity representation. Crucially, we embed batch-wise negative prototypes into the diffusion process and train under objectives that combine cross-entropy reconstruction with a cosine-alignment regularizer. This formulation not only sharpens the model’s ability to discriminate between plausible and implausible future events but also preserves the generative strengths of diffusion dynamics. Comprehensive experiments conducted across four widely-used TKG benchmarks demonstrate that NADEx consistently achieves state-of-the-art performance.

Acknowledgments

We thank the anonymous reviewers for their valuable discussion and feedback. This work was supported by the National Natural Science Foundation of China (U22B2061), and the National Key R&D Program of China (2022YFB4300603).

Limitation

One limitation of this work lies in its reliance on four widely used TKG benchmarks (ICEWS14, ICEWS05–15, ICEWS18, and GDELT), although standard, represent a relatively narrow spectrum of event data. These datasets are dominated by political and international relations events, with limited diversity in domains such as science, economics, or natural disasters. As a result, the model’s generalizability to other types of temporal knowledge graphs remains untested. In addition, while the proposed negative-aware diffusion framework introduces a cosine-alignment regularizer to incorporate negative prototypes, the construction of negatives is limited to batch-level sampling. This strategy may not fully capture semantically hard negatives or long-tail entities, potentially leading to overly coarse approximations of implausible events.

Ethics Statement

This study adheres to established ethical standards. We rely solely on datasets that were previously collected and annotated by earlier research, and our use of these resources entails no new data gathering, no interaction with human subjects, and no processing of personally identifiable or otherwise private information. The models and results are presented for scientific analysis and benchmarking and are not intended for surveillance, deception, or any form of harm. In all cases, we respect data-use terms and privacy safeguards and ensure that our work does not endanger the rights, safety, or dignity of any individual or group.

References

- Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, and 1 others. 2022. Gpt-neox-20b: An open-source autoregressive language model. In *Proceedings of Big-Science Episode# 5–Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.
- Elizabeth Boschee, Jennifer Lautenschlager, Sean O’Brien, Steve Shellman, James Starz, and Michael Ward. 2015. [ICEWS Coded Event Data](#).
- Borui Cai, Yong Xiang, Longxiang Gao, He Zhang, Yunfeng Li, and Jianxin Li. 2023. Temporal knowledge graph completion: a survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 6545–6553.
- Yuxiang Cai, Qiao Liu, Yanglei Gan, Changlin Li, Xueyi Liu, Run Lin, Da Luo, and Jiaye Yang. 2024. [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, Bangkok, Thailand. Association for Computational Linguistics.
- Yukun Cao, Lisheng Wang, and Luobin Huang. 2025. Dpcl-diff: Temporal knowledge graph reasoning based on graph node diffusion model with dual-domain periodic contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 14806–14814.
- Kai Chen, Xin Song, Ye Wang, Liqun Gao, Aiping Li, Xiaojuan Zhao, Bin Zhou, and Yalong Xie. 2025a. Llm-dr: A novel llm-aided diffusion model for rule generation on temporal knowledge graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 11481–11489.
- Tingxuan Chen, Jun Long, Zidong Wang, Shuai Luo, Jincui Huang, and Liu Yang. 2024a. Thcn: A hawkes process based temporal causal convolutional network for extrapolation reasoning in temporal knowledge graphs. *IEEE Transactions on Knowledge and Data Engineering*.
- Tingxuan Chen, Liu Yang, Zidong Wang, Shuai Luo, and Jun Long. 2025b. Enhancing extrapolation reasoning on temporal knowledge graphs with logic rules and queries. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Wei Chen, Huaiyu Wan, Yuting Wu, Shuyuan Zhao, Jiayaqi Cheng, Yuxin Li, and Youfang Lin. 2024b. 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. 2025c. Cogntke: A cognitive temporal knowledge extrapolation framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 14815–14823.
- Wei-Lin Chiang, Zhuohan Li, Ziqing Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, and 1 others. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.

- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794.
- Hao Dong, Zhiyuan Ning, Pengyang Wang, Ziyue Qiao, Pengfei Wang, Yuanchun Zhou, and Yanjie Fu. 2023. Adaptive path-memory network for temporal knowledge graph reasoning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 2086–2094.
- Alberto Garcia-Duran, 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*, pages 4816–4821.
- Rishab Goel, Seyed Mehran Kazemi, Marcus Brubaker, and Pascal Poupart. 2020. Diachronic embedding for temporal knowledge graph completion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 3988–3995.
- Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and Lingpeng Kong. 2022. Diffuseq: Sequence to sequence text generation with diffusion models. In *The Eleventh International Conference on Learning Representations*.
- Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and Lingpeng Kong. 2023. Diffuseq-v2: Bridging discrete and continuous text spaces for accelerated seq2seq diffusion models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9868–9875.
- Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. 2021. [Explainable subgraph reasoning for forecasting on temporal knowledge graphs](#). In *International Conference on Learning Representations*.
- Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and S Yu Philip. 2021. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE transactions on neural networks and learning systems*, 33(2):494–514.
- 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.
- Seyed Mehran Kazemi and David Poole. 2018. Simple embedding for link prediction in knowledge graphs. *Advances in neural information processing systems*, 31.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. 2020. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*.
- Julien Leblay and Melisachew Wudage Chekol. 2018. Deriving validity time in knowledge graph. In *Companion proceedings of the the web conference 2018*, pages 1771–1776.
- Dong-Ho Lee, Kian Ahrabian, Woojeong Jin, Fred Morstatter, and Jay Pujara. 2023. Temporal knowledge graph forecasting without knowledge using in-context learning. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 544–557.
- Kalev Leetaru and Philip A. Schrodt. 2013. [Gdelt: Global data on events, location, and tone](#). *ISA Annual Convention*, 2:1–49.
- Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. 2022a. Diffusion-lm improves controllable text generation. *Advances in neural information processing systems*, 35:4328–4343.
- Yujia Li, Shiliang Sun, and Jing Zhao. 2022b. Tirgn: Time-guided recurrent graph network with local-global historical patterns for temporal knowledge graph reasoning. In *IJCAI*, pages 2152–2158.
- Zixuan Li, Saiping Guan, Xiaolong Jin, Weihua Peng, Yajuan Lyu, Yong Zhu, Long Bai, Wei Li, Jiafeng Guo, and Xueqi Cheng. 2022c. Complex evolutionary pattern learning for temporal knowledge graph reasoning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 290–296.
- Zixuan Li, Zhongni Hou, Saiping Guan, Xiaolong Jin, Weihua Peng, Long Bai, Yajuan Lyu, Wei Li, Jiafeng Guo, and Xueqi Cheng. 2022d. Hismatch: Historical structure matching based temporal knowledge graph reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7328–7338.
- 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, SIGIR '21*, page 408–417, New York, NY, USA. Association for Computing Machinery.
- Ke Liang, Lingyuan Meng, Meng Liu, Yue Liu, Wenxuan Tu, Siwei Wang, Sihang Zhou, Xinwang Liu, Fuchun Sun, and Kunlun He. 2024. A survey of knowledge graph reasoning on graph types: Static, dynamic, and multi-modal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

- 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.
- Haohu Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and Mark D Plumbley. 2023a. Audioldm: Text-to-audio generation with latent diffusion models. In *International Conference on Machine Learning*, pages 21450–21474. PMLR.
- Kangzheng Liu, Feng Zhao, Guandong Xu, Xianzhi Wang, and Hai Jin. 2023b. 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.
- Yushan Liu, Yunpu Ma, Marcel Hildebrandt, Mitchell Joblin, and Volker Tresp. 2022. Tlogic: Temporal logical rules for explainable link forecasting on temporal knowledge graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 4120–4127.
- Ruilin Luo, Tianle Gu, Haoling Li, Junzhe Li, Zicheng Lin, Jiayi Li, and Yujiu Yang. 2024. Chain of history: Learning and forecasting with llms for temporal knowledge graph completion. *CoRR*.
- Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2022. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning*, pages 16784–16804. PMLR.
- Renning Pang, Yao Liu, Yanglei Gan, Tingting Dai, Yashen Wang, Xiaojun Shi, Tian Lan, and Qiao Liu. 2025. Improving temporal knowledge graph reasoning with hierarchical semantic-aware contrastive learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 376–394. Springer.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Diffusionner: Boundary diffusion for named entity recognition. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3875–3890.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. pmlr.
- 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*, pages 8306–8319.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothee Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models.
- Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. 2017. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. In *international conference on machine learning*, pages 3462–3471. PMLR.
- Jiapu Wang, Sun Kai, Linhao Luo, Wei Wei, Yongli Hu, Alan Wee-Chung Liew, Shirui Pan, and Baocai Yin. 2024. Large language models-guided dynamic adaptation for temporal knowledge graph reasoning. *Advances in Neural Information Processing Systems*, 37:8384–8410.
- Wenjie Wang, Yiyang Xu, Fuli Feng, Xinyu Lin, Xiangan He, and Tat-Seng Chua. 2023. Diffusion recommender model. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 832–841.
- Yi Xu, Junjie Ou, Hui Xu, and Luoyi Fu. 2023. Temporal knowledge graph reasoning with historical contrastive learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 4765–4773.
- Bishan Yang, Scott Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding entities and relations for learning and inference in knowledge bases. In *Proceedings of the International Conference on Learning Representations (ICLR) 2015*.
- Zhengyi Yang, Jiancan Wu, Zhicai Wang, Xiang Wang, Yancheng Yuan, and Xiangan He. 2023. Generate what you prefer: Reshaping sequential recommendation via guided diffusion. *Advances in Neural Information Processing Systems*, 36:24247–24261.
- Mengqi Zhang, Yuwei Xia, Qiang Liu, Shu Wu, and Liang Wang. 2023. Learning long-and short-term representations for temporal knowledge graph reasoning. In *Proceedings of the ACM web conference 2023*, pages 2412–2422.
- Shengzhe Zhang, Wei Wei, Rikui Huang, Wenfeng Xie, and Danyang Chen. 2024. Modeling historical relevant and local frequency context for representation-based temporal knowledge graph forecasting. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7675–7686.
- Shangfei Zheng, Hongzhi Yin, Tong Chen, Quoc Viet Hung Nguyen, Wei Chen, and Lei Zhao. 2023. Dream: Adaptive reinforcement learning based on attention mechanism for temporal knowledge graph

reasoning. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*, pages 1578–1588.

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.

A Additional Related Work

A.1 Diffusion Models on Discrete Data

Diffusion models (DMs) (Sohl-Dickstein et al., 2015) have emerged as a unifying generative paradigm, delivering state-of-the-art sample quality in image generation (Dhariwal and Nichol, 2021; Nichol et al., 2022) and audio (Kong et al., 2020; Liu et al., 2023a). While the original formulations operate in continuous Euclidean space, several lines of work now transpose the framework to discrete symbol sequences. Diffusion-LM (Li et al., 2022a) projects word tokens into continuous embeddings and performing score matching within that continuous latent space, enabling effective non-autoregressive and controlled text generation, whereas DiffuSeq (Gong et al., 2022, 2023) adopts a discrete corruption process aligned with time indices and employs a transformer-based denoiser, enabling the generation of coherent sequences with arbitrary length. Beyond sequence generation, diffusion has also been adapted for structured prediction. DiffusionNER (Shen et al., 2023), for instance, formulates named-entity recognition as a boundary denoising task, gradually refining noisy span boundaries into consistent entity predictions.

More recently, diffusion-based frameworks have been introduced into domain-specific applications that involve symbolic or sequential interactions. In recommendation systems, DiffRec (Wang et al., 2023) and DreamRec (Yang et al., 2023) leverage diffusion to model user–item interaction sequences. By injecting noise into observed behaviors and learning to reverse this process, these models uncover preference distributions and capture fine-grained uncertainty in recommendation outcomes. Such extensions highlight the versatility of diffusion paradigms as they are not only effective in dense continuous domains but also adaptable to discrete symbolic reasoning tasks, where noise-driven generative processes can reveal hidden structure and probabilistic dependencies.

B Preliminary

Definition 3. Diffusion Models for Discrete Data.

Diffusion models (DMs) are probabilistic generative frameworks comprising two coupled Markov chains: a **forward diffusion** that gradually corrupts data with noise, and a **reverse denoising** that learns to recover the original samples. Although originally formulated for continuous domains, recent works have extended DMs to discrete sequences.

We summarize the key components of discrete-data diffusion below.

Forward diffusion. Given a discrete sequence w , we first embed it into continuous space via,

$$\mathbf{x}_0 \sim q(\mathbf{x}_0 | w) = \mathcal{N}(\mathbf{x}_0; \text{Emb}(w), \beta_0 \mathbf{I}), \quad (18)$$

where $\text{Emb}(\cdot) \in \mathbb{R}^d$ is an mapping function that projects each word into vector. $\beta_0 > 0$ controls the initial noise level. Thereafter, for $t \in \{0, 1, \dots, T\}$, the latent states evolve according to :

$$\begin{aligned} q(\mathbf{x}_t | \mathbf{x}_{t-1}) &= \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_{t-1}, (1 - \bar{\alpha}_t) \mathbf{I}) \\ &= \sqrt{\bar{\alpha}_t} \mathbf{x}_{t-1} + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \end{aligned} \quad (19)$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(0, 1)$ is a random Gaussian noise, and $\bar{\alpha}_t = \prod_{t'=1}^t \alpha_{t'} \in (0, 1)$ governs the noise schedule.

Reverse denoising. Starting from $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the reverse chain is parameterized by a neural network f_θ that predicts the Gaussian transition,

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t); \Sigma_\theta(\mathbf{x}_t, t)), \quad (20)$$

where μ_θ and Σ_θ are learnable functions of (\mathbf{x}_t, t) . By iteratively sampling from p_θ , one can recover an estimate $\hat{\mathbf{x}}_0$.

Discrete rounding and training objective. To map $\hat{\mathbf{x}}_0$ back to tokens, we apply a soft-max based decoder,

$$p_\theta(\mathbf{w} | \hat{\mathbf{x}}_0) = \text{Softmax}(\hat{\mathbf{x}}_0). \quad (21)$$

The end-to-end loss combines a continuous reconstruction term and a discrete rounding term,

$$\begin{aligned} \mathcal{L}_{\text{simple}}(\mathbf{w}) &= \mathbb{E}_{q_\phi(\mathbf{x}_{0:T} | \mathbf{w})} \left[\sum_{t=2}^T [\|\mathbf{x}_0 - f_\theta(\mathbf{x}_t, t)\|^2] \right] + \\ &\mathbb{E}_{q_\phi(\mathbf{x}_{0:1} | \mathbf{w})} [\|\text{Emb}(\mathbf{w}) - f_\theta(\mathbf{x}_1, 1)\|^2 - \log p_\theta(\mathbf{w} | \mathbf{x}_0)]. \end{aligned} \quad (22)$$

Where the first term trains the model to denoise from $t = 2$ to T , reducing continuous approximation errors, while the second term aligns the $t = 1$ denoised output with the original embedding and encourages accurate rounding to the discrete sequence.

Table 7: The statistics of the datasets. $|E|$ and $|R|$ denote the number of unique entities and event types, respectively.

Datasets	$ E $	$ R $	Train	Valid	Test	Unseen Ratio
ICEWS14	6,869	230	74,845	8,514	7,371	58.43%
ICEWS18	23,033	256	373,018	45,995	49,545	55.69%
ICEWS05-15	10,094	251	368,868	46,302	46,159	39.82%
GDELT	7,691	240	1,734,399	238,765	305,241	43.72%

C Experimental Setup

C.1 Dataset Statistics

To ensure consistency and comparability, we adhere to the chronological splitting protocol (80% training, 10% validation, and 10% testing) established in prior studies by (Li et al., 2021; Cai et al., 2024).

C.2 Baselines

Static Baselines:

- DistMult (Yang et al., 2015), employs a bilinear scoring function, modeling triple plausibility via a diagonal relation matrix that captures pairwise interactions between subject and object embeddings.
- ConvE (Dettmers et al., 2018), applies 2D convolution over reshaped entity and relation embeddings, followed by a projection layer to learn richer feature interactions for link prediction.
- RotatE (Sun et al., 2019), represents relations as complex-valued rotations in the embedding space, enabling the model to naturally encode and infer diverse relational patterns such as symmetry and inversion.

Interpolation Baselines:

- TTransE (Leblay and Chekol, 2018), explicitly models temporal dynamics by embedding entities and relations within a continuous time framework, using translation operations along the time dimension to capture their evolution.
- TA-DistMult (Garcia-Duran et al., 2018), extends the DistMult scoring function with time-aware embeddings, allowing the model to adapt relation parameters according to temporal context.
- DE-Simple (Goel et al., 2020), employs diachronic embeddings that parameterize entity

and relation representations as functions of time, hence modeling progressive change across different timestamps.

Extrapolation Baselines:

- RE-NET (Jin et al., 2020), integrates recurrent neural architectures with graph convolution to capture the sequential evolution of entities and predict future links.
- Re-GCN (Li et al., 2021), employs a Recurrent Evolutionary GCN that recurrently updates entity and relation embeddings at each timestamp by propagating temporal signals through the KG.
- CyGNet (Zhu et al., 2021), models cyclical temporal patterns, enabling the framework to learn periodic behaviors and extrapolate yet-unseen links.
- TITer (Sun et al., 2021), applies hierarchical transformations to entity embeddings, iteratively tracking their evolution to anticipate future TKG states.
- CEN (Li et al., 2022c), uses length-aware convolutional filters to extract multi-scale evolutionary patterns, with an online training strategy to handle temporal variability.
- TiRGN (Li et al., 2022b), leverages recurrent graph networks to encode dynamic relational structures, improving inference of unseen facts.
- HisMatch (Li et al., 2022d), aligns historical entity representations via a matching mechanism, enhancing prediction accuracy for future interactions.
- RETIA (Liu et al., 2023b), constructs a twin hyper-relation subgraph and evolutionarily aggregates adjacent entity and relation features for enriched message passing.
- CENET (Xu et al., 2023), incorporates contrastive learning objectives to strengthen dynamic representation learning within temporal graphs.
- CRAFT (Zhang et al., 2024), utilizes hypergraph convolutional aggregators to fuse coarse- and fine-grained historical contexts, yielding robust extrapolation.

- THCN (Chen et al., 2024a), introduces temporal causal convolutional networks grounded in Hawkes processes to distinguish the relative importance of concurrent facts.
- DiffuTKG (Cai et al., 2024), reframes TKG reasoning as a denoising diffusion process over entity embeddings to generate plausible future links.
- LogiQ (Chen et al., 2025b), augments diffusion-based generation with logical constraints, improving both accuracy and interpretability.
- CognTKE (Chen et al., 2025c), integrates cognitively inspired symbolic priors into embedding learning, enabling more transparent and reliable extrapolation.

D Experimental Analysis

D.1 Sensitivity Analysis

In this section, we evaluate NADEX’s sensitivity to three key hyperparameters on the ICEWS14 and ICEWS18 benchmarks in terms of MRR and Hit@1, including the impact of temperature coefficient τ , different lengths of event sequence and different scale of noise in the forward process. Our findings are summarized in Figure 4.

Impact of Temperature Coefficient. The temperature coefficient τ in our log-likelihood ranking loss controls the sharpness of the softmax distribution over candidate entities. From Figure 4(a) on ICEWS14, we observe that both MRR and Hit@1 improve as τ increases from 0.1 to 0.5. Beyond this point, performance plateaus or slightly dips. In contrast, on the more temporally complex ICEWS18 (Figure 4(d)), performance initially drops at $\tau = 0.3$, but then steadily rises across the range $\tau = 0.5$ to $\tau = 0.9$, with optimal MRR and Hit@1 achieved at $\tau = 0.9$.

Impact of Event Sequence Length. The length of the historical event sequence notably influences NADEX’s predictive power, as shown in Figures 4(b) and (e). On ICEWS14, increasing the input window from 4 to 8 events yields a substantial jump, indicating that the model benefits from additional context. Further extending the sequence to 16, 32, and 64 events provides diminishing but still positive returns. The plateau around 32 to 64 suggests extra history adds limited new information on ICEWS14. ICEWS18 exhibits a similar pattern, with steady MRR and Hit@1 increases up to 64

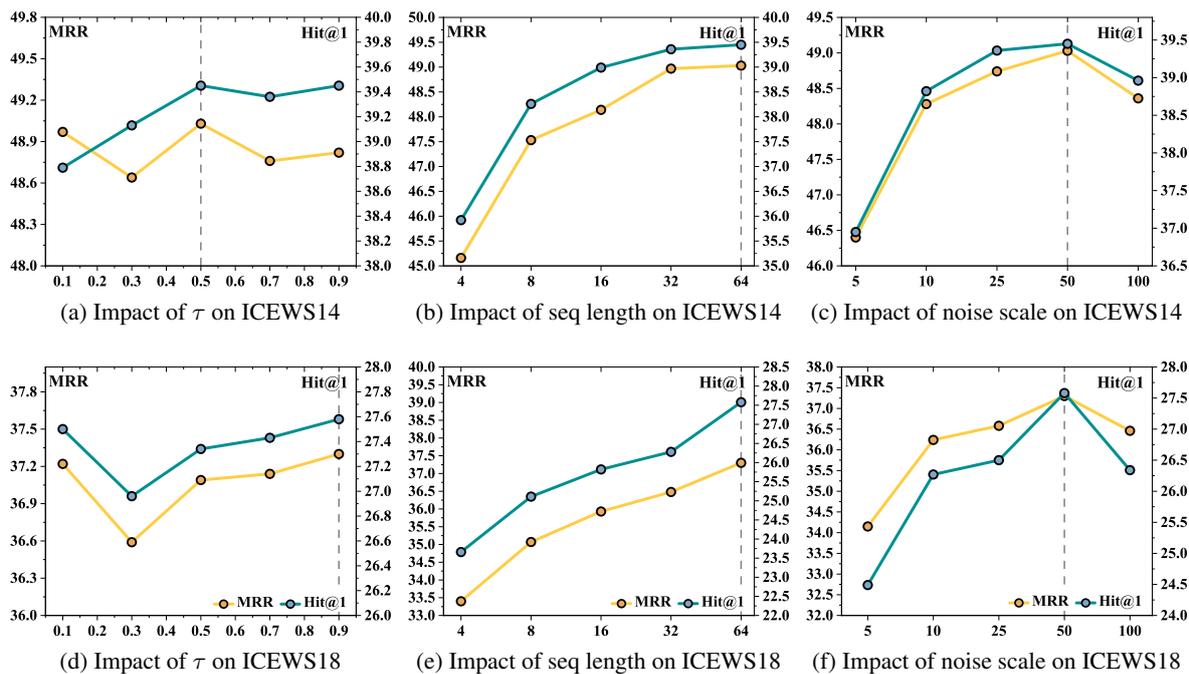


Figure 4: Sensitivity analysis on ICEWS14 and ICEWS18.

events, and a relative acceleration in improvement beyond 32, reflecting its more complex temporal dynamics. It is worth-noting that we restrict our analysis to a maximum sequence length of 64 to balance predictive gains against the computational overhead on a single NVIDIA A100 (80 GB) GPU.

Impact of Noise Scale. We further investigate the influence of the maximum noise scale in the forward diffusion process on NADEX’s capacity to model and denoise intricate temporal dynamics. Figures 4(c) and (f) report MRR and Hit@1 as the perturbation magnitude applied to the target embedding is varied. When the scale is small (5–10), the injected noise is insufficient to exploit the generative power of diffusion, resulting in sub-optimal accuracy. Performance improves markedly at a moderate scale of 25 and reaches its optimum at 50 on both ICEWS14 and ICEWS18, indicating that a balanced level of stochasticity best enriches the denoising signal. Conversely, very large perturbations drive the latent representation far from the data manifold; the reverse process must then traverse a longer, noisier trajectory, weakening gradient signals and making accurate reconstruction substantially more difficult.