# Event-Driven Cascade Updates in Knowledge Graphs

Bo Ni
bo.ni@vanderbilt.edu
Vanderbilt University
Nashville, Tennessee, USA

Tyler Derr
tyler.derr@vanderbilt.edu
Vanderbilt University
Nashville, Tennessee, USA

## Abstract

Knowledge Graphs provide structured representations of real-world entities and relationships, but they are often incomplete or outdated due to delayed updates. To address this challenge, we introduce the task of consistency-aware knowledge graph updating, where the goal is to predict structural changes induced by a triggering event in order to preserve the logical consistency in the knowledge graph. We further propose a text-augmented encoder-decoder framework that uses LLM-generated descriptions to provide semantic signals that guide a relational graph convolutional network to capture both direct and cascading effects. To support this task, we also construct ICEWS14-Event, a novel benchmark dataset specifically designed to evaluate consistency-preserving KG updates. Preliminary experimental results demonstrate that our method is able to outperform state-of-the-art baselines in correctly updating knowledge graphs.

## Keywords

Knowledge Graphs, Graph Reasoning, Cascading Updates.

**ACM Reference Format:**
Bo Ni and Tyler Derr. 2018. Event-Driven Cascade Updates in Knowledge Graphs. *ACM/IMS J. Data Sci.* 37, 4, Article 111 (August 2018), 3 pages. https://doi.org/XXXXXXX.XXXXXXX

## 1 Introduction

Knowledge Graphs (KGs) are widely used in applications such as question answering [4], recommendation systems [2], and information retrieval [1]. However, most existing methods assume that KGs are static, factual, and up-to-date [6]. Thus, ensuring the reliability and consistency of KGs emerges as a critical challenge [5].

While prior work such as Tang et al. [8] explore KG updates using textual signals, but their method depends heavily on the presence of rich natural language input and is limited to modeling shallow, local changes. This makes it difficult to capture cascading multi-hop relational effects triggered by complex events.

To address this gap, we propose a text-augmented encoder-decoder framework that combines event representations generated by large language models with a relational graph encoder to propagate structural updates. Our method is designed to reason over both local and multi-hop relational consequences of events, enabling more consistent KG updates.

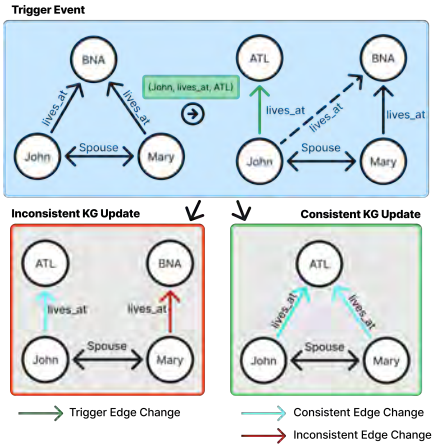


**Figure 1: An illustration of *Consistent* KG Updates**

We also introduce ICEWS14-Event, a new benchmark dataset constructed by aligning event data with KG snapshots and applying consistency-preserving relational rules. This dataset captures both explicit and cascading effects from real-world events, allowing for a more holistic evaluation of consistency-aware KG models.

Our contributions are summarized as follows:

- We identify and formulate the problem of consistency-aware KG updates with cascading relational effects.
- We propose a novel graph reasoning framework that encodes event semantics and predicts multi-hop structural updates.
- We construct the ICEWS14-Event dataset, which includes aligned trigger events and rule-based consistency updates.

## 2 Problem Definition

Let $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ be a knowledge graph with entities $\mathcal{E}$, relations $\mathcal{R}$, and factual triplets $\mathcal{T} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$. Let $e$ be an external event that induces structural changes in the graph.

*Definition 2.1 (Event-based KG Update).* Given a KG $\mathcal{G}$ and a trigger event $e$, the task is to predict a revised set of triplets $\mathcal{T}'$ such that the updated graph $\mathcal{G}' = (\mathcal{E}, \mathcal{R}, \mathcal{T}')$ reflects the new factual state. This involves learning a function $f(e, \mathcal{G}) \rightarrow \Delta\mathcal{T}$, where $\Delta\mathcal{T}$ denotes the added or removed triplets.

*Remark* 1 (Consistency Objective). Unlike general KG forecasting [**?** ], our focus is on maintaining logical consistency immediately after a triggering event. For example, updating the spouse's *lives_in* relation after relocation (Figure 1) preserves logical coherence.

*Definition 2.2 (Cascading KG Updates).* A cascading update refers to the recursive propagation of predicted changes $\mathcal{T}'$ throughout the graph to preserve consistency across multi-hop dependencies.

Authors' Contact Information: Bo Ni, bo.ni@vanderbilt.edu, Vanderbilt University, Nashville, Tennessee, USA; Tyler Derr, tyler.derr@vanderbilt.edu, Vanderbilt University, Nashville, Tennessee, USA.


ACM 2831-3194/2018/8-ART111
https://doi.org/XXXXXXX.XXXXXXX

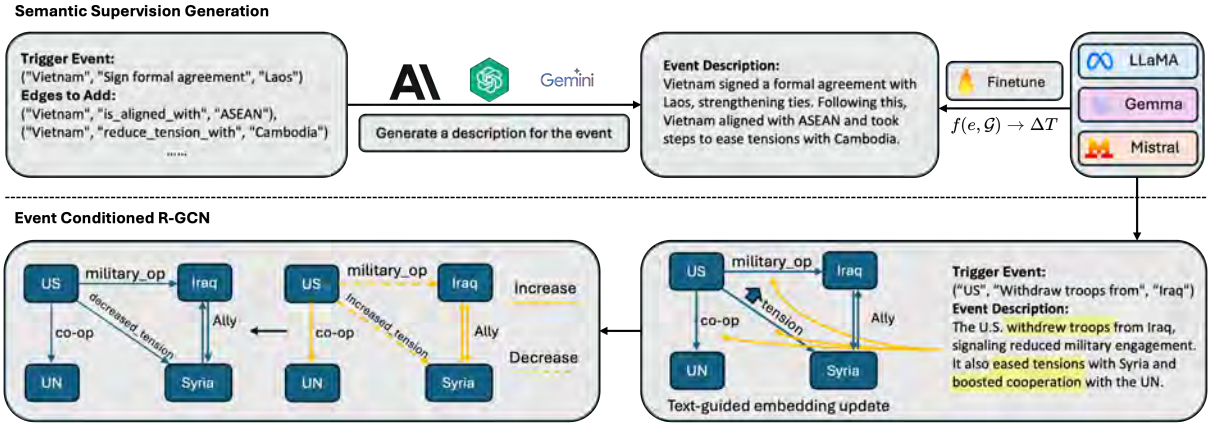


**Figure 2: Overview of the proposed CascadeKG.**

## 3 ICEWS14-Event

Due to the dearth of research on consistency-aware knowledge graph updates, we construct a new dataset, ICEWS14-Event, to facilitate evaluation in this setting.

The dataset aligns real-world events with KG snapshots and applies a set of hand-crafted consistency rules to propagate relational changes beyond explicitly stated facts. These rules are applied recursively to simulate realistic, multi-hop cascades triggered by a single event, capturing the logical dependencies between relations. We recognize this dataset as the first benchmark focused on evaluating logical consistency in event-driven KG update.

## 4 Methods

**Framework Overview.** As illustrated in Figure 2, our approach introduces a text-augmented encoder-decoder framework built on R-GCN [7] to maintain consistency during KG updates. Let $g$ be a predefined transformation function that converts structured inputs into natural language prompts. Given an external event $e$, we construct a prompt $\mathcal{P} = g(e, \Delta\mathcal{T})$ describing the structural changes $\Delta T$ in the knowledge graph, and use a commercial LLM $M_c$ to generate natural language context $x = M_c(\mathcal{P})$. This supervision guides the fine-tuning of an open-source model $M_s$ via LoRA [3], optimizing:

$$\mathcal{L}_{\text{gen}}(\theta) = \mathbb{E}_{(e,\Delta\mathcal{T},x)} \left[-\log P_\theta(x \mid e, \Delta\mathcal{T})\right]$$

The generated descriptions are encoded into semantic embeddings $z_x = f_{\text{enc}}(x)$, projected via $\tilde{z}_x = W_t z_x$, and broadcast into the graph as initial node states $h_v^{(0)} = h_v^{\text{raw}} + \tilde{z}_x$.

To capture the impact of $e$, we apply our text-conditioned R-GCN to both $\mathcal{G}$ and $\mathcal{G} + e$, and concatenate the resulting embeddings:

$$z_v = [z_v^{\text{orig}} \,\|\, z_v^{\text{mod}}]$$

These fused embeddings are then decoded to predict triplets $\mathcal{T}'$.

**Cascading Updates.** KG updates can propagate indefinitely through relational dependencies. To prevent runaway inference and control error propagation, we introduce a discount factor $\gamma \in (0, 1]$. At each hop $k$, inferred triplets are weighted by $\gamma^k$, downscaling the influence of more distant inferences.

| | KG Consistency 1-Hop | | Cascading Updates 3-Hop | |
|---|---|---|---|---|
| Baseline | F1 | MRR | F1 | MRR |
| R-GCN | 95.5 | 59.7 | 96.6 | 59.0 |
| R-GAT [8] | 94.8 | 42.0 | **96.8** | 23.6 |
| DistMult | 80.7 | 20.1 | 92.5 | 13.3 |
| TransE | 0 | 17.5 | 0 | 16.8 |
| CascadeKG | **97.3** | **61.7** | **96.8** | **61.1** |

**Table 1: Preliminary Results on ICEWS14-Event.**

## 5 Experiments

We report our preliminary results in Table 1 under two settings. **KG Consistency** evaluates the model's ability to predict changes within the 1-hop neighborhood of the event, while **Cascading Updates** measures performance on multi-hop relational effects triggered by a single event. Following prior work [8], we report standard metrics including F1 score and Mean Reciprocal Rank (MRR). For baselines, we include both embedding-based baselines (DistMult, TransE) and state-of-the-art encoder-based models (R-GCN, R-GAT). Our preliminary results suggest that the proposed CascadeKG achieves improved consistency across both settings.

## 6 Conclusion and Future Work

In this work, we identified a new problem of consistency-aware knowledge graph updates and proposed CascadeKG, a text-augmented framework for modeling event-induced structural changes. Our approach leverages LLM-generated descriptions and a text-conditioned encoder to improve multi-hop relational reasoning. To support this task, we introduce ICEWS14-Event, the first benchmark dataset specifically designed to evaluate consistency and cascading effects in event-driven KG updates. We are currently expanding our experiments to larger event-driven datasets, specifically extending the setting to YAGO3-10. Moreover, we plan to design improved prompts to help LLMs capture update-relevant meta paths more effectively. Finally, while our current cascading mechanism relies on a discount factor heuristic, future work will explore uncertainty-aware halting strategies for more principled propagation control.

## References


[1] Jinheon Baek, Alham Fikri Aji, Jens Lehmann, and Sung Ju Hwang. 2023. Direct Fact Retrieval from Knowledge Graphs without Entity Linking. arXiv:2305.12416 [cs.IR] https://arxiv.org/abs/2305.12416

[2] Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. 2020. A Survey on Knowledge Graph-Based Recommender Systems. arXiv:2003.00911 [cs.IR] https://arxiv.org/abs/2003.00911

[3] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*. https://openreview.net/forum?id=nZeVKeeFYf9

[4] Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2024. Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning. arXiv:2310.01061 [cs.CL] https://arxiv.org/abs/2310.01061

[5] Bo Ni, Zheyuan Liu, Leyao Wang, Yongjia Lei, Yuying Zhao, Xueqi Cheng, Qingkai Zeng, Luna Dong, Yinglong Xia, Krishnaram Kenthapadi, Ryan Rossi, Franck Dernoncourt, Md Mehrab Tanjim, Nesreen Ahmed, Xiaorui Liu, Wenqi Fan, Erik Blasch, Yu Wang, Meng Jiang, and Tyler Derr. 2025. Towards Trustworthy Retrieval Augmented Generation for Large Language Models: A Survey. arXiv:2502.06872 [cs.CL] https://arxiv.org/abs/2502.06872

[6] Bo Ni, Yu Wang, Lu Cheng, Erik Blasch, and Tyler Derr. 2024. Towards Trustworthy Knowledge Graph Reasoning: An Uncertainty Aware Perspective. arXiv:2410.08985 [cs.AI] https://arxiv.org/abs/2410.08985

[7] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *The semantic web: 15th international conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, proceedings 15*. Springer, 593–607.

[8] Jizhi Tang, Yansong Feng, and Dongyan Zhao. 2019. Learning to Update Knowledge Graphs by Reading News. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 2632–2641. doi:10.18653/v1/D19-1265