Computers and Artificial Intelligence; Vol.3, No.6; 2026
E-ISSN: 3006-3728 P-ISSN: 3006-371X
DOI: https://doi.org/10.70267/cai.v3n6.133140
Published by: Zeus Press

TAR-KU: Temporal-Aware Risk-Constrained Knowledge Updating for High-Recall Political Text Correction

Zhu Sun*

College of Computer Science and Technolog, National University of Defense Technology, Changsha, China

**Corresponding author: Zhu Sun.*

Abstract

Political texts contain time-sensitive facts about people, offices, institutions, places, and official terminology. A knowledge base that is merely expanded with newly retrieved material may improve coverage, but it can also introduce entity confusion, outdated facts, conflicting evidence, and unjustified edits. This paper presents TAR-KU, a temporal-aware risk-constrained knowledge updating method for high-recall political text correction. TAR-KU separates knowledge admission from correction authorization. Incremental materials are first converted into candidate records, then checked through entity-relation alignment, temporal validity modeling, source-evidence aggregation, conflict governance, and version snapshots. Only records that satisfy the query-time authorization conditions are allowed to support local edits. This design preserves historical knowledge while making newly verified facts available to downstream correction. Experiments on 688 political-text samples, together with an independent 128-sample source set and a 100-sample clean control set, compare TAR-KU with retrieval, memory, parameter-editing, snapshot, and prompt-update baselines. TAR-KU achieves 83.19% correction precision, 87.59% correction recall, and 85.34% F1, obtaining the highest recall and F1 among the compared methods, at an average latency of 4.07 s. Updating the knowledge state increases recall from 85.63% to 87.59% and correct edits from 566 to 579. The results show that TAR-KU improves dynamic-fact coverage while preserving an explicit boundary between evidence availability and automatic modification.

Keywords

political text correction, temporal knowledge, knowledge updating, retrieval-augmented generation, entity-relation alignment, risk-constrained editing

1. Introduction

Political text correction differs from ordinary spelling correction because correctness depends not only on linguistic well-formedness but also on the identity, role, institution, location, and validity time of the mentioned facts. A sentence can be fluent while assigning a current office to a former official, linking an institution to the wrong region, or using a historical title in a current report. These errors are particularly difficult for a high-recall proofreading system: missing a real error may allow it to enter publication, whereas a suspicious suggestion can still be checked by a human editor.

© 2026 The Author(s). Authors retain copyright in their articles.

Retrieval-augmented generation (RAG) provides a practical way to supply changing information without repeatedly retraining a language model [1]. In-context retrieval, active retrieval, and self-reflective retrieval further study when evidence should be acquired and how it should influence generation [2–4]. However, semantic relevance is not equivalent to temporal validity or entity identity. Two documents may both mention the same office while describing different people and different periods. Increasing the number of retrieved passages can therefore improve evidence coverage and simultaneously increase the chance of irrelevant or contradictory candidates.

Knowledge editing offers another route by changing model parameters or attaching an external memory. ROME and MEMIT perform targeted and mass factual editing [5,6], while SERAC and later lifelong editing methods use external memories, routing, or specialized editing modules [7–10]. These approaches make new facts accessible to the model, but a parameter state or memory item does not by itself preserve all historical versions, identify the exact effective interval, or authorize a particular character-level replacement. Political text correction needs a mechanism that connects knowledge maintenance with local edit decisions without treating the latest record as a universal replacement.

This paper proposes TAR-KU (Temporal-Aware Risk-Constrained Knowledge Updating). Its central principle is to separate three decisions: whether a newly extracted record can enter a maintained knowledge state, which version is applicable to a query, and whether the applicable evidence is strong enough to authorize a local correction. The method combines incremental discovery, entity-relation alignment, temporal constraints, evidence aggregation, conflict detection, version snapshots, and edit authorization. The downstream correction model is kept as a separate execution component; TAR-KU supplies it with evidence and permissions rather than replacing its entire generation process.

The contributions are threefold. First, TAR-KU represents dynamic political knowledge as versioned records with explicit entities, relations, validity conditions, sources, and review states. Second, it introduces a query-time authorization layer that prevents every newly stored record from becoming an unconditional editing instruction. Third, it evaluates the method using both quality and cost measures, including a clean control set, update-state comparison, component ablation, knowledge-scale variation, and knowledge-prompt budget variation.

2. Related Work

2.1 Retrieval and Temporal Knowledge

RAG couples document retrieval with conditional generation and is useful when the knowledge source changes more frequently than model parameters [1]. Active retrieval and Self-RAG make retrieval selective or reflective [2,3], while graph-based retrieval organizes entities and relations to improve multi-hop evidence access [11]. Temporal knowledge-graph research updates representations in streaming settings and models historical or evolving relations [12–15]. These studies provide important building blocks, but political text correction requires a further mapping from a temporally valid relation to a bounded span and an allowed edit.

2.2 Knowledge Editing and External Memory

Parameter editing can make factual changes available inside the model. ROME locates and edits factual associations, and MEMIT extends this idea to multiple facts [5,6]. WISE, WilKE, and regularized editing methods address lifelong editing, memory organization, or collateral damage [8–10]. SERAC stores edits externally and routes inputs according to applicability [7]. In-context editing methods such as IKE use examples rather than direct weight modification [16]. Their common limitation in this application is that knowledge access and edit authorization are usually not identical operations. TAR-KU therefore treats them as separate stages and evaluates both coverage and modification risk.

2.3 Temporal Evidence and Editing Reliability

TimeR4 retrieves temporal knowledge-graph evidence using time-aware constraints [17]; MRAG uses modular retrieval for time-sensitive question answering [18]. These approaches motivate the use of temporal

© 2026 The Author(s). Authors retain copyright in their articles.

information when selecting evidence, whereas TAR-KU additionally connects evidence to local editing decisions.

RippleEdits investigates effects of factual edits on related knowledge [19], and Hsueh et al. examine the limitations of knowledge editing [20]. Further studies consider historical information and temporal granularity in knowledge editing [21,22]. Continuous prompting and external editing memory offer alternative approaches to integrating new facts [23,24]. Together, these works motivate explicit consideration of history, update effects, and applicability.

3. Problem Formulation

Let a political text be represented as a character sequence x. A correction candidate is a local edit e=(s,t,op), where s is the original span, t is the target span, and op is insertion, deletion, or replacement. A knowledge record is represented as k=(h,r,o,τ,src,v), where h and o are entities, r is a relation, τ is a validity interval, src is source evidence, and v is its version state.

The update problem is not simply to maximize the number of stored records. For an incoming record, the system must determine whether the entities and relation are aligned, whether its validity interval is compatible with existing records, whether evidence is sufficiently independent and reliable, and whether the record should be active, pending, rejected, or retained only for historical reference. For a correction request, the system must additionally determine whether the record applies to the text time and context and whether it supports a unique local edit. A record may therefore be stored without being authorized for automatic correction.

4. The TAR-KU Method

4.1 Overview

TAR-KU follows the sequence: incremental discovery → structured extraction → entity-relation alignment → temporal and evidence validation → conflict governance → version publication → query-time filtering → edit authorization. The update side changes the knowledge resource; the query side reads a selected version and decides whether evidence can support an edit. This separation avoids updating the entire resource during every correction request and prevents a newly observed statement from directly overwriting historical knowledge.

4.2 Incremental Discovery and Structured Records

New material is first processed as an untrusted candidate rather than as executable knowledge. The system extracts entities, relation types, attributes, source identifiers, publication dates, possible effective dates, and textual evidence. Duplicate or near-duplicate reports are grouped so that repeated syndication is not incorrectly counted as independent support. Candidate records retain the original evidence span, making later review and rollback possible.

4.3 Entity-relation Alignment

Alignment checks whether the candidate refers to the same person, office, institution, or place as an existing record. Surface similarity alone is insufficient for similar names, translated names, abbreviations, and titles. TAR-KU uses entity attributes and relation direction to constrain matching. For person–office relations, it can query both “who holds this office?” and “which office does this person hold?” For country–capital relations, the reverse relation is likewise checked. Region, institution, and role restrictions are preserved rather than discarded during extraction.

4.4 Temporal Validity and Conflict Governance

Each fact is associated with an effective interval or time condition. Historical and current facts can coexist, while query-time filtering selects the record compatible with the text date. Conflict detection is applied to records sharing the same subject, relation, and overlapping validity interval with mutually exclusive objects; legitimate multiple-value or concurrent relations are not automatically treated as conflicts. Sources are

© 2026 The Author(s). Authors retain copyright in their articles.

aggregated with their provenance and review state. Unresolved conflicts remain pending instead of being converted into automatic edits.

4.5 Version Snapshots and Edit Authorization

Published knowledge is organized into version snapshots. A snapshot records active, historical, pending, and rejected states, as well as the source and review trail. During correction, TAR-KU filters the selected snapshot using the text time, entity context, relation compatibility, and evidence strength. A candidate is authorized only when the target entity and relation are compatible, the original span is fully identifiable, the edit direction is unambiguous, and no unresolved mutually exclusive conflict remains. Authorized knowledge is converted into a local edit and passed to the downstream candidate-fusion module. Knowledge support improves the decision score but does not grant permission to rewrite an entire sentence.

4.6 Complexity and Operational Behavior

Incremental indexing avoids reprocessing unchanged documents, but alignment and conflict checks may still grow with the number of new entities and records. Grouping records by entity and relation and filtering by time reduce the practical search space. The method also introduces storage for evidence, snapshots, and review states and adds query-time validation. Consequently, TAR-KU is designed to trade additional verification cost for more controlled coverage rather than to minimize latency under all conditions.

5. Experimental Setup

Experiments use Qwen2.5-7B-Instruct with the same LoRA adapter and downstream correction configuration. The main test set contains 688 political-text samples. An independent T2 source set contains 128 samples, and a 100-sample clean control set is used to inspect unnecessary modifications. The news case base contains 6,715 records. The initial knowledge state covers reviewed records available in July–August 2026; the updated state additionally includes reviewed records through September 20, 2026. Retrieval uses top-4 evidence, and the per-sample scheduling budget is approximately 5 s. All comparisons use the same random seed (42), data, model, and evaluation rules.

Correction precision measures the proportion of predicted edits that match the reference edit, while correction recall measures the proportion of reference edits correctly recovered. F1 is their harmonic mean. Average latency is measured at the sample level. Baselines represent distinct knowledge-use paths: static RAG, fixed snapshot retrieval, IKE-inspired in-context updating, MEMIT-inspired editing memory, SERAC-inspired routing, temporal snapshot switching, conflict-aware incremental RAG, and local prompt updating. These are mechanism-adapted baselines rather than claims of exact reproduction of every original system.

6. Results and Discussion

6.1 Comparison with Baseline Methods

Table 1: Comparison of Baseline Methods

Method	Precision (%)	Recall (%)	F1 (%)	Latency (s)
MEMIT-inspired	83.18	81.54	82.35	2.42
Temporal snapshot switching	83.59	81.69	82.63	2.37
Conflict-aware incremental RAG	84.42	82.00	83.19	2.42
Local prompt updating	85.15	81.54	83.31	2.54
SERAC-inspired routing	81.25	76.70	78.91	2.49
Static RAG	80.34	77.91	79.11	2.43
Fixed snapshot retrieval	80.50	78.06	79.26	2.46
IKE-inspired updating	86.31	81.09	83.62	2.60
TAR-KU	83.19	87.59	85.34	4.07

TAR-KU obtains the highest recall and F1. Relative to static RAG, recall increases by 9.68 percentage points and F1 by 6.23 points. Relative to the strongest external baseline in recall, conflict-aware incremental RAG, recall increases by 5.59 points. Relative to the strongest external baseline in F1, IKE-inspired updating, F1 increases by 1.72 points. TAR-KU does not obtain the highest precision: IKE-inspired updating reaches

© 2026 The Author(s). Authors retain copyright in their articles.

86.31%, compared with 83.19% for TAR-KU. Its average latency is also higher than the 2.37–2.60 s range of the baselines.

The pattern is consistent with the position of TAR-KU in the correction chain. Static retrieval and prompt or memory updating mainly improve how new facts enter the request. TAR-KU additionally generates relation-based candidates and checks identity, temporal conditions, evidence, and conflicts before authorizing local edits. This expands the set of recoverable dynamic-fact errors, but the additional candidate and validation stages create extra processing cost. The result is therefore a high-recall, evidence-controlled operating point rather than universal superiority on every metric.

6.2 Effect of Knowledge-state Updating

Table 2:Effect of Knowledge-State Updating

Knowledge state	Precision (%)	Recall (%)	F1 (%)	Correct edits	Latency (s)
Initial	82.63	85.63	84.10	566	4.12
Updated	83.19	87.59	85.34	579	4.07

After reviewed incremental knowledge is added, correct edits increase from 566 to 579, recall increases by 1.96 points, precision increases by 0.56 points, and F1 increases by 1.24 points. Latency changes from 4.12 to 4.07 s, a small difference that should not be interpreted as a significant acceleration. The result indicates that the updated records supply evidence for errors that were previously unsupported, while the same temporal and authorization checks constrain their use.

6.3 Ablation Study

Table 3:Ablation Results for TAR-KU Components

Configuration	Precision (%)	Recall (%)	F1 (%)	Latency (s)
Full configuration	83.19	87.59	85.34	4.07
Without relation candidates	84.01	81.85	82.91	4.12
Weak validation fusion	82.55	87.29	84.85	4.08
Without entity approximation	83.31	87.59	85.40	4.10
Without clean-text protection	82.83	87.59	85.15	4.05
Without RAG reranking	82.76	87.14	84.89	4.06

Removing relation candidates causes the clearest coverage loss: recall decreases from 87.59% to 81.85%, while precision rises by 0.82 points. This shows that structured relation candidates contribute errors that are difficult to find through surface similarity alone. Weak validation fusion and removal of RAG reranking produce smaller decreases in recall or F1, suggesting that evidence selection contributes to, but does not solely determine, the result. Removing entity approximation leaves recall unchanged and slightly raises F1 to 85.40%; therefore, this component does not provide a net gain under the current test conditions. Latency remains within 4.05–4.12 s across variants.

6.4 Knowledge Scale and Prompt Budget

Table 4: Effect of Incremental Knowledge Scale

Incremental knowledge (%)	Precision (%)	Recall (%)	Latency (s)
0	72.06	82.37	4.28
25	72.77	83.81	4.26
50	73.83	86.85	4.19
75	72.86	87.28	4.19
100	72.78	88.87	4.15

Increasing the incremental knowledge scale from 0% to 100% raises recall from 82.37% to 88.87%. Precision improves until 50% and then returns to approximately 72.8%, indicating that additional knowledge mainly expands coverage and does not guarantee a simultaneous precision gain. The small latency variation is not monotonic and remains within 4.15–4.28 s.

Adding three knowledge items improves recall from 87.86% to 88.87% and precision from 71.95% to 72.78%. Larger budgets keep recall in the narrow range of 88.44–88.87% and do not produce a stable

© 2026 The Author(s). Authors retain copyright in their articles.

additional gain. This suggests that a small set of high-relevance evidence can be sufficient for many requests, while more evidence may introduce competition in the prompt.

Table 5:Effect of Knowledge Prompt Budget

Knowledge prompt budget	Precision (%)	Recall (%)	Latency (s)
0	71.95	87.86	4.17
3	72.78	88.87	4.11
6	72.69	88.87	4.14
9	72.51	88.44	4.12
12	72.78	88.87	4.13

7. Limitations

The experiments cover two knowledge states rather than a long sequence of continuous updates. They do not fully quantify concurrent service capacity, peak memory, or long-term rollback reliability. The evaluation is also focused on political-text correction and uses mechanism-adapted baselines; results should not be generalized to all domains or interpreted as exact reproductions of the original baseline implementations. Further work should expand authoritative sources, improve entity disambiguation and relation extraction, evaluate forward-only temporal updates, and report stage-level costs and multiple random seeds.

8. Conclusion

This paper presented TAR-KU, a temporal-aware risk-constrained knowledge updating method for high-recall political text correction. By separating knowledge admission, version selection, and edit authorization, TAR-KU preserves historical facts while making newly reviewed evidence available to local correction. On the evaluated political-text benchmark, it achieves the highest recall and F1 among the compared methods, while incurring additional latency and not maximizing precision. Knowledge-state, ablation, scale, and prompt-budget experiments consistently show that broader evidence access mainly improves error coverage; evidence validation and authorization remain necessary to control unjustified edits. TAR-KU therefore provides a practical mechanism for dynamic knowledge maintenance in human-assisted political text proofreading.

References

[1] Lewis P, Perez E, Piktus A, et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. NeurIPS, 2020: 9459–9474.

[2] Ram O, Levine Y, Dalmedigos I, et al. In-Context Retrieval-Augmented Language Models. TACL, 2023, 11: 1316–1331.

[3] Asai A, Wu Z, Wang Y, et al. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. ICLR, 2024.

[4] Jiang Z, Xu F, Gao L, et al. Active Retrieval Augmented Generation. EMNLP, 2023: 7969–7992.

[5] Meng K, Bau D, Andonian A, Belinkov Y. Locating and Editing Factual Associations in GPT. NeurIPS, 2022: 17359–17372.

[6] Meng K, Sharma A S, Andonian A, et al. Mass-Editing Memory in a Transformer. ICLR, 2023.

[7] Mitchell E, Lin C, Bosselut A, et al. Memory-Based Model Editing at Scale. ICML, 2022, 162: 15817–15831.

[8] Wang P, Li Z, Zhang N, et al. WISE: Rethinking the Knowledge Memory for Lifelong Model Editing of Large Language Models. NeurIPS, 2024, 37: 53764–53797.

[9] Hu C, Cao P, Chen Y, et al. WilKE: Wise-Layer Knowledge Editor for Lifelong Knowledge Editing. Findings of ACL, 2024: 3476–3503.

© 2026 The Author(s). Authors retain copyright in their articles.

[10] Gu J C, Xu H X, Ma J Y, et al. Model Editing Harms General Abilities of Large Language Models: Regularization to the Rescue. EMNLP, 2024: 16801–16819.

[11] Peng B, Zhu Y, Liu Y, et al. Graph Retrieval-Augmented Generation: A Survey. ACM TOIS, 2025, 44(2): 1–52.

[12] Zhang J, Shao J, Cui B. StreamE: Learning to Update Representations for Temporal Knowledge Graphs in Streaming Scenarios. SIGIR, 2023: 622–631.

[13] Chen Z, Xu C, Su F, et al. Temporal Extrapolation and Knowledge Transfer for Lifelong Temporal Knowledge Graph Reasoning. Findings of EMNLP, 2023: 6736–6746.

[14] Zhang F, Zhang Z, Zhuang F, et al. Temporal Knowledge Graph Reasoning With Dynamic Memory Enhancement. IEEE TKDE, 2024, 36(11): 7115–7128.

[15] Xin H, Chen L. KartGPS: Knowledge Base Update with Temporal Graph Pattern-based Semantic Rules. ICDE, 2024: 5075–5087.

[16] Zheng C, Li L, Dong Q, et al. Can We Edit Factual Knowledge by In-Context Learning? EMNLP, 2023: 4862–4876.

[17] Qian X, Zhang Y, Zhao Y, et al. TimeR4: Time-aware Retrieval-Augmented Large Language Models for Temporal Knowledge Graph Question Answering. EMNLP, 2024: 6942-6952. DOI: 10.18653/v1/2024.emnlp-main.394.

[18] Zhang S, Xue Y, Zhang Y, et al. MRAG: A Modular Retrieval Framework for Time-Sensitive Question Answering. Findings of EMNLP, 2025: 3080-3118. DOI: 10.18653/v1/2025.findings-emnlp.167.

[19] Cohen R, Biran E, Yoran O, Globerson A, Geva M. Evaluating the Ripple Effects of Knowledge Editing in Language Models. TACL, 2024, 12: 283-298. DOI: 10.1162/tacl_a_00644.

[20] Hsueh C H, Huang P K M, Lin T H, et al. Editing the Mind of Giants: An In-Depth Exploration of Pitfalls of Knowledge Editing in Large Language Models. Findings of EMNLP, 2024: 9417-9429. DOI: 10.18653/v1/2024.findings-emnlp.550.

[21] Yin X, Jiang J, Yang L, Wan X. History Matters: Temporal Knowledge Editing in Large Language Model. AAAI, 2024, 38(17): 19413-19421. DOI: 10.1609/aaai.v38i17.29912.

[22] Zhao S, Pang N, Tan Z, et al. Multi-granularity Temporal Knowledge Editing over Large Language Models. AAAI, 2026, 40(34): 28742-28750. DOI: 10.1609/aaai.v40i34.40107.

[23] Chen Q, Zhang T, He X, et al. Lifelong Knowledge Editing for LLMs with Retrieval-Augmented Continuous Prompt Learning. EMNLP, 2024: 13565-13580. DOI: 10.18653/v1/2024.emnlp-main.751.

[24] Qiao S, Liu X, Na S H. COMEM: In-Context Retrieval-Augmented Mass-Editing Memory in Large Language Models. Findings of NAACL, 2024: 2333-2347. DOI: 10.18653/v1/2024.findings-naacl.151.

Funding

This research received no external funding.

Conflicts of Interest

The author(s) declare no conflicts of interest.

Acknowledgment

This paper is an output of the science project.

© 2026 The Author(s). Authors retain copyright in their articles.

Copyrights

Copyright for this article is retained by the author(s), with first publication rights granted to the journal. This is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (http://creativecommons.org/licenses/by/4.0/).

© 2026 The Author(s). Authors retain copyright in their articles.