
000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 EMERGE: A BENCHMARK FOR UPDATING KNOWLEDGE GRAPHS WITH EMERGING TEXTUAL KNOWLEDGE

Anonymous authors

Paper under double-blind review

ABSTRACT

Knowledge Graphs (KGs) are structured knowledge repositories containing entities and relations between them. In this paper, we study the problem of automatically updating KGs over time in response to evolving knowledge in unstructured textual sources. Addressing this problem requires identifying a wide range of update operations based on the state of an existing KG at a given time *and the information extracted from text*. This contrasts with traditional information extraction pipelines, which extract knowledge from text independently of the current state of a KG. To address this challenge, we propose a method for construction of a dataset consisting of Wikidata KG snapshots over time and Wikipedia passages paired with the corresponding edit operations that they induce in a particular KG snapshot. *We obtain these pairs by aligning annotated hyperlinked entity mentions in each Wikipedia passage with the corresponding entities involved in the updated Wikidata triples.* We verify, using LLMs with human validation, that these textual passages contain the knowledge needed to support the associated KG edits. The resulting dataset comprises 233K Wikipedia passages aligned with a total of 1.45 million KG edits over 7 different yearly snapshots of Wikidata from 2019 to 2025. Our experimental results highlight key challenges in updating KG snapshots based on emerging textual knowledge, particularly in integrating knowledge expressed in text with the existing KG structure. These findings position the dataset as a valuable benchmark for future research. We will publicly release our dataset and model implementations.¹

1 INTRODUCTION

Knowledge graphs (KGs) play a crucial role in applications such as question answering (Wang et al., 2024; Dong et al., 2025), recommender systems (Zhang et al., 2024; Wang et al., 2025), information retrieval (Reinanda et al., 2020), fact-checking (Kim et al., 2023; Hao & Wu, 2025), and healthcare prediction (Jiang et al., 2025), among others (Zou, 2020). Furthermore, KGs provide structured, queryable world knowledge that increasingly complements *large language models (LLMs)* (Pan et al., 2024; Cai et al., 2025). This integration has been used to reduce LLM hallucinations (Agrawal et al., 2024; Lavrinovics et al., 2025), improve fine-tuning (Chen et al., 2025; Ma et al., 2025), enhance planning (Chen et al., 2024; Petruzzellis et al., 2025), support complex reasoning (Sun et al., 2024; Luo et al., 2025), and provide reliable knowledge augmentation (Han et al., 2024; Li et al., 2025c). However, as world knowledge evolves, KGs must also be updated to remain reliable (Polleres et al., 2023; Hofer et al., 2024; Li et al., 2025b). Yet, existing temporal KG benchmarks (Liang et al., 2024; Alam et al., 2024) model only internal temporal dynamics and do not capture how KGs should be updated in response to new world knowledge emerging in external textual sources. In addition, current textual information extraction (IE) datasets and models (Zhao et al., 2024b; Xu et al., 2024) do not link extracted facts to the concrete KG updates they should induce.

To address these limitations, we introduce EMERGE, a novel, automatically constructed benchmark that aligns emerging textual knowledge with the concrete updates it induces in a KG. Concretely, EMERGE links evolving changes in the Wikidata KG (Vrandečić & Krötzsch, 2014) with

¹Code and dataset will be released upon acceptance. The test set is included in the supplementary material.

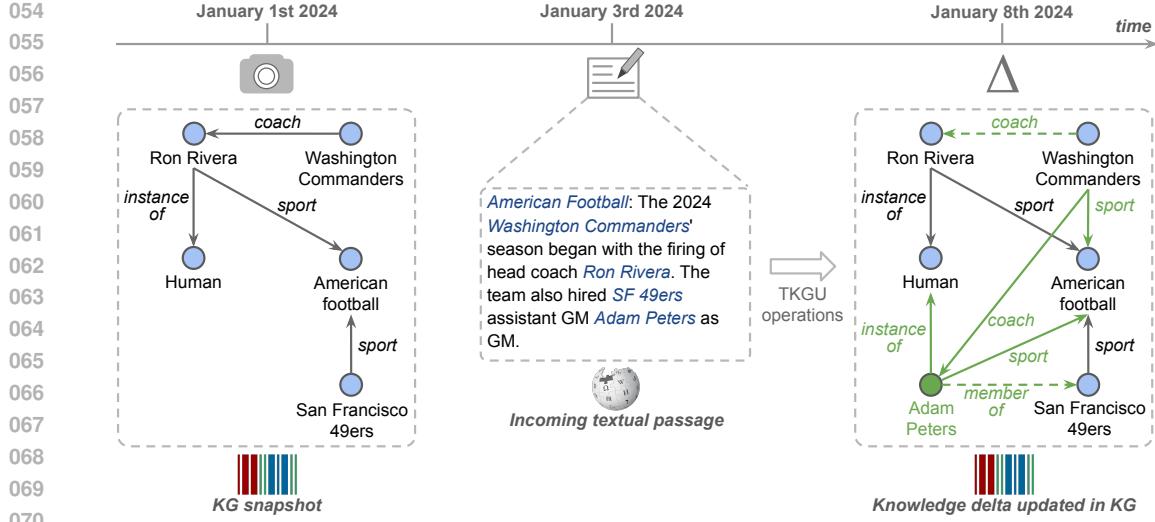


Figure 1: Illustration of one instance in EMERGE. The reference KG *snapshot* of January 1st 2024 is updated with new, *emerging knowledge* contained in the incoming *textual passage* from January 3rd 2024. The *updated KG* involves not only creation of new relations (solid green arrows), but also generation of new entities (green circle) and deprecation of relations (dashed green arrows).

corresponding textual passages from Wikipedia that reflect these KG updates over time. This alignment enables evaluating both (i) how well models integrate new textual knowledge into a KG and (ii) how temporally evolving KG structures (i.e., KG schema) affect this integration. Furthermore, EMERGE is incrementally extensible through an automatic pipeline that continuously incorporates new knowledge from Wikipedia and Wikidata. Figure 1 illustrates an example in which a *KG snapshot* from January 1st, 2024 is updated based on *emerging textual evidence* from January 3rd, 2024. To construct such instances, EMERGE aligns weekly *knowledge deltas* in Wikidata with the corresponding textual changes in Wikipedia.

Furthermore, EMERGE differs from existing mainstream IE datasets, which primarily focus on extracting triples from text, either via manual annotations or by linking text to a static KG. As a result, these existing benchmarks cannot evaluate the broader set of operations needed to update a KG to reflect textual knowledge. Such operations require creation new entities, linking them to existing ones, and deprecation of outdated facts (see Figure 1 for an example). To capture these requirements, EMERGE defines five text-driven KG updating (TKGU) operations (see Section 3) and is, to our knowledge, the first dataset to support all of them (see a detailed comparison with existing benchmarks in Appendix A.1). Our benchmarking further shows that state-of-the-art IE models fall short in supporting the full range of operations required to update KGs with new knowledge (see Table 1). Furthermore, these models rely solely on knowledge expressed in text and remain unaware of how that knowledge is structured within a KG. As a result, the extracted triples, though semantically valid, often fail to align with the KG schema and structure.

In summary, the contributions of this paper are as follows:

- We formalize and study the problem of maintaining KGs from emerging textual knowledge, defining it through a set of fundamental text-driven KG updating (TKGU) operations.
- EMERGE, a novel dataset that maps emerging knowledge in textual passages to corresponding updates in temporally evolving KG snapshots.
- A publicly available pipeline for extending EMERGE with new KG snapshots, enabling the evaluation of models on continuously evolving knowledge.
- Experimental results and analysis on EMERGE using two state-of-the-art IE architectures.

108
109
110
111
112
113

Table 1: Comparison of state-of-the-art information extraction models by the type of extracted knowledge: (1) existing KG triples (*X-Triples*), (2) new triples with existing KG entities (*E-Triples*), (3) new triples with emerging entities (*EE-Triples*), (4) new triples linking emerging entities to the rest of the KG (*EE-KG-Triples*), and (5) deprecated triples (*D-Triples*). The *KG Link* column indicates whether extracted triples are linked to a KG.

Model	KG Link	Supported textual knowledge type extraction				
		<i>X-Triples</i>	<i>E-Triples</i>	<i>EE-Triples</i>	<i>EE-KG-Triples</i>	<i>D-Triples</i>
REBEL (2021)	✗	✓	✓	✓	✗	✗
GenIE (2022)	✓	✓	✓	✗	✗	✗
KnowGL (2023)	✓	✓	✓	✓	✗	✗
GCD (2023)	✓	✓	✓	✗	✗	✗
ReLiK cIE (2024)	✓	✓	✓	✗	✗	✗
ReLiK RE (2024)	✗	✓	✓	✓	✗	✗
EDC (2024)	✗	✓	✓	✓	✗	✗
ATG (2024)	✗	✓	✓	✓	✗	✗
CodeKGC (2024)	✗	✓	✓	✓	✗	✗

125
126
127
128

2 RELATED WORK

129
130
131

Below we describe the most relevant related work directions. Additionally, we provide an extensive related work section and a comparison table (Table 4) in Appendix A.

KG completion and refinement. Research on KG completion (KGC) (Shen et al., 2022) and refinement (Paulheim, 2016; Subagdja et al., 2024) has produced many datasets aimed at predicting missing relations between entities. Early work introduced WN18 and FB15k Bordes et al. (2013), derived from WordNet (Miller, 1995) and Freebase (Bollacker et al., 2008), followed by improved variants such as WN18RR and FB15k-237 (Toutanova & Chen, 2015; Dettmers et al., 2018) addressing redundancy and data leakage. Larger and more recent datasets include Wikidata5M Wang et al. (2021), along with Wiki/NELL-One (Xiong et al., 2018), FB15K-237N (Lv et al., 2022), CoDEX (Safavi & Koutra, 2020), YAGO3-10 (Mahdisoltani et al., 2014), and LiterallyWikidata (Gesese et al., 2021). While these datasets evaluate models on predicting new edges within the KG, they remain restricted to the KG internal structure. Our objective instead is to support KG updates driven by the information originating in external unstructured textual sources. This distinction also separates our work from temporal KG completion (TKGC) datasets such as GDELT (Leetaru & Schrot, 2013), ICEWS14/05-15 (Garcia-Duran et al., 2018), Wikidata12k (Dasgupta et al., 2018), Wikidata-big (Lacroix et al., 2020), ICEWS18 (Jin et al., 2020), and more recently TGB and TGB 2.0 (Huang et al., 2024; Gastinger et al., 2024), among others (Liang et al., 2024; Alam et al., 2024). While these benchmarks capture internal temporal evolution of facts, they do not model how KGs should be kept up to date with knowledge emerging in external textual sources. EMERGE fills this gap by aligning such external textual evidence with the concrete KG updates it induces, enabling the study of models that keep KGs updated as world knowledge evolves.

149
150
151
152
153
154
155
156
157
158
159
160
161

Information extraction (IE). To evaluate the ability of models to extract structured knowledge, researchers have developed IE datasets by annotating entity relations. MUC-7 (Chinchor & Marsh, 1998) introduced three relation types, with later datasets expanding in size, relation diversity, or both. Notable examples include CoNLL04 (Roth & Yih, 2004), ACE 2005 (Walker et al., 2006), ERE (Aguilar et al., 2014; Song et al., 2015), BC5CDR (Li et al., 2016), TACRED (Zhang et al., 2017), SciERC (Luan et al., 2018), SemEval-2010 (Hendrickx et al., 2010), SemEval-2017 (Augenstein et al., 2017), DWIE (Zaporozets et al., 2021) and BioRED (Luo et al., 2022), among others. Other datasets, such as NYT (Riedel et al., 2010), explicitly linked KG triples to textual snippets using distant supervision. Similarly, but on a larger scale, Gabrilovich et al. (2013) introduced FACC1 by aligning ClueWeb12 documents with Freebase entity mention annotations. In parallel, the TAC-KBP challenges (Ji et al., 2010; TAC-KBP, 2022) (2009 – 2020) produced proprietary manually annotated datasets for knowledge base population tasks such as slot filling and entity linking. More recently, these resources have been extended with a variety of datasets that map textual knowledge to KG triples, such as WebNLG (Gardent et al., 2017), KELM (Agarwal et al., 2021), FewRel

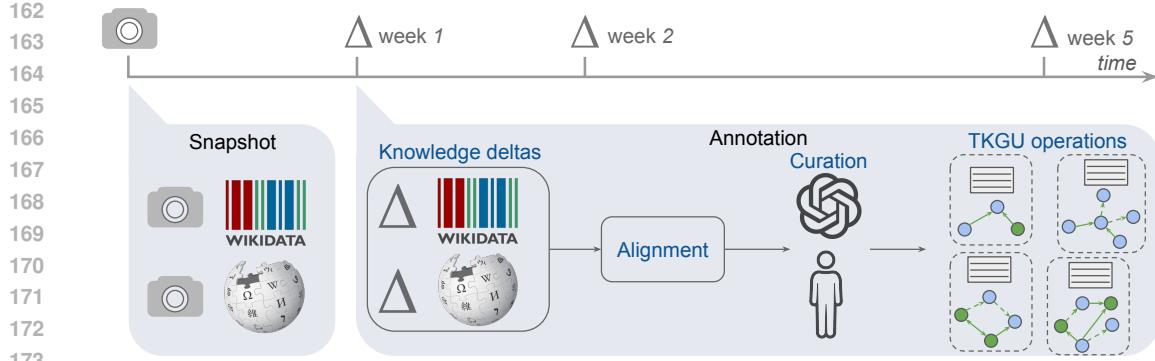


Figure 2: Illustration of EMERGE creation pipeline. First, weekly *knowledge deltas* (Δ) are extracted by identifying changes in Wikipedia passages and Wikidata KG relative to a fixed *snapshot*. In the *Alignment* step, these KG and textual deltas are connected. During *Curation*, an LLM discards KG updates not supported by aligned textual changes, a process verified with manual annotations on a subsample of alignments. The result is high-quality text–KG update pairs, as in Figure 1, where multiple *TKGU operations* (Section 3) update the KG with emerging textual knowledge.

(Han et al., 2018), DocRED (Yao et al., 2019), Wiki/GEO-NRE (Distiawan et al., 2019), BioRel (Xing et al., 2020), T-REX (Elsahar et al., 2019) and REBEL (Cabot & Navigli, 2021). While these datasets connect textual knowledge to KG triples and literals, they do not account for the operations required to update a KG as new information emerges in text. Our work addresses this gap by linking new textual knowledge to the specific update operations (see Section 3) on a KG snapshot. Methodologically, existing state-of-the-art IE methods (see Table 1) provide a natural starting point for tackling TKGU operations, as they extract structured knowledge from text and already cover some of these operations. However, they remain largely oblivious to the existing KG structure and require extensions to integrate emerging textual content into the KG, such as deprecating outdated triples, adding new entities, and enforcing structural consistency based on how entities and relations are used within the KG.

3 PROBLEM DEFINITION

We define the problem of *text-driven knowledge graph updating* (TKGU) as determining the necessary edits to a KG at a particular point in time, given a textual passage. More formally, we define a KG *snapshot* at some point time t as a tuple $G_t = (V_t, R_t, T_t)$ where V_t is a set of entities, R_t is a set of relation types, and T_t is a set of *triples* of the form (s, p, o) where $s, o \in V_t$ are the subject and object, and $p \in R_t$ is the relation between them. Given a textual passage $d_{t'}$ created at some point in time $t' > t$, the task consists in generating a set of TKGU operations defined as follows:

Emerging triples (E-Triples). Addition of triples that are not present in the KG but involve entities that already exist in it; that is, $(s, p, o) \notin G_t$ and $s \in V_t \wedge o \in V_t$. For example, in Figure 1, the added triple (*Washington Commanders*, *sport*, *American Football*) involves the entities *Washington Commanders* and *American Football*, both of which already exist in the KG.

Emerging entities and triples (EE-Triples). Addition of triples that do not exist in the KG and involve a subject entity, object entity, or both that are also absent. That is, $(s, p, o) \notin G_t$ and $s \notin V_t \vee o \notin V_t$. For example, in Figure 1, the added triple (*Washington Commanders*, *coach*, *Adam Peters*) introduces the entity *Adam Peters*, which is not yet in the KG.

Emerging entities to KG triples (EE-KG-Triples). Addition of new triples where exactly one of the subject or object entities is mentioned in a passage $d_{t'}$, while the other already exists in the KG and is not explicitly mentioned in the passage. These triples evaluate the ability of the models to integrate newly emerging entities by linking them to existing ones in the KG. For example, in Figure 1, the added triple (*Adam Peters*, *instance of*, *Human*) links the emerging entity *Adam Peters* to the existing entity *Human*, even though this relation is not explicitly stated in the passage.

216 **Deprecated triples (D-Triples).** Deprecation of triples already existing in a KG based on emerging
217 evidence in textual passage. For example, in Figure 1, the triples (*Adam Peters, member of,*
218 *San Francisco 49ers*) and (*Washington Commanders, coach, Ron Rivera*) are deprecated based on
219 updated information in the passage.

220 **Existing triples (X-Triples).** Detection of triples already existing in the KG that are supported by
221 textual passage, i.e., $(s, p, o) \in G_t$. This operation evaluates the ability of models to recognize
222 existing knowledge. For example, in Figure 1, the triple (*San Francisco 49ers, sport, American*
223 *football*) is both supported by the passage and already present in the original KG snapshot.

225 Table 1 compares existing IE architectures based on the types of TKGU operations (see above) they
226 are able to extract. While many models can extract triples involving existing entities in a KG (*X-*
227 *Triples* and *E-Triples*), most struggle to identify triples with emerging entities (*EE-Triples*) and none
228 of them supports linking them to the rest of the KG (*EE-KG-Triples*). Furthermore, some methods
229 only partially integrate newly extracted knowledge, as they do not link the extracted triples to the KG
230 (see *KG* column). For example, *relation extraction* models such as REBEL (Cabot & Navigli, 2021)
231 and ReLiK RE (Orlando et al., 2024) are able to extract new triples but do not link their entities
232 and relations to the KG; other models such as EDC (Zhang & Soh, 2024), link only relations but
233 not entities. Finally, existing IE methods, to the best of our knowledge, are not designed to identify
234 triples that should be deprecated based on emerging textual knowledge (*D-Triples*).

235 4 OUR DATASET

238 We introduce EMERGE, a large-scale dataset that, unlike existing benchmarks, supports all the
239 TKGU operations defined in Section 3.

241 4.1 DATA COLLECTION

243 We construct a dataset consisting of 7 Wikidata yearly snapshots taken on January 1st at 00:00 GMT
244 from 2019 to 2025. We expect that these snapshots will enable to evaluate the drift in temporal
245 performance of models pre-trained at different time points. To evaluate the ability of the models to
246 update KG with emerging knowledge, we generate cumulative weekly *deltas* (up to 5 weeks) for
247 each snapshot (see Figure 2). Each delta represents a time window and includes textual passages
248 along with the corresponding KG updates occurring during that period. Below, we describe in more
249 detail the main steps in the EMERGE dataset creation pipeline.

250 **Wikipedia and Wikidata dumps.** We begin by downloading the historical revision logs from the
251 Wikipedia and Wikidata dumps available at <https://dumps.wikimedia.org/>. These logs
252 provide complete access to the revision history of Wikipedia and Wikidata, enabling fine-grained
253 tracking of temporal changes. Using this level of granularity, we are able to construct EMERGE using
254 *any number of arbitrarily defined KG snapshots and delta windows, with temporal precision*
255 *down to the second*. This capability sets EMERGE apart from existing datasets designed to evaluate
256 model performance on evolving KG knowledge (Boschee et al., 2015; Dasgupta et al., 2018; Lacroix
257 et al., 2020), which are typically derived from a single KG snapshot and rely only on temporal
258 attributes associated with edges. While such datasets are valuable for predicting the emergence of new
259 facts over time, they do not allow the evaluation of how structural changes in the KG across different
260 snapshots affect model performance. Moreover, because we have access to the full revision history
261 of Wikipedia pages, we can evaluate models on all the newly introduced textual content within any
262 chosen temporal delta. This allows us to assess, for instance, how varying the size of delta windows
263 influences model performance. It also contrasts with related datasets using Wikipedia (Lewis et al.,
264 2020; Jang et al., 2022a; Onoe et al., 2023; Zhao et al., 2024a), which are based on only one or a
265 small number of manually downloaded Wikipedia snapshots, thereby limiting temporal flexibility.

266 **Snapshot generation.** Given a list of desired snapshot timestamps, we process Wikipedia and
267 Wikidata history revisions to obtain the following components for each timestamp t : (1) a Wikidata
268 KG snapshot G_t corresponding to t , (2) a dictionary of entities present in Wikipedia at t , along with
269 their corresponding textual descriptions, and (3) a dictionary of relation types present in Wikidata at
270 t with definitions. In line with the Wikidata5M dataset (Wang et al., 2021), we restrict the Wikidata
271 KG to include only entities that are present in Wikipedia.

270 **KG deltas generation.** For each snapshot, we generate deltas in weekly increments, spanning up to
271 5 weeks. Each delta represents the difference between two KG snapshots, denoted as $G_{t+\Delta} - G_t$,
272 where Δ represents the delta window. Each of the resulting deltas involve KG triple operations
273 outlined in Section 3. Concretely, *X-Triples* exist in G_t and $G_{t+\Delta}$, *E-Triples* contain new emerging
274 relations in $G_{t+\Delta}$ between entities already existing in G_t , and *EE-Triples* and *EE-KG-Triples* consist
275 of emerging relations between entities where subject or object do not exist in G_t , and is introduced in
276 $G_{t+\Delta}$. Finally, to obtain *D-Triples*, besides including the removed edges, we match Wikidata triple
277 qualifiers (see Appendix J) that explicitly indicate knowledge **deprecation** within the delta interval.
278 We mark these triples as **deprecated** rather than removing them, since the underlying fact does not
279 change but expires within the delta interval.

280 **Aligning KG deltas with text.** For each delta in a given snapshot t , we retrieve the newly introduced
281 Wikipedia passages within the temporal window corresponding to that delta. Following the approach
282 of Cabot & Navigli (2021); Elsahar et al. (2019), we then *align* these passages with triples in each
283 of the KG deltas by matching the annotated hyperlinked entity mentions in each of the passages to
284 the corresponding entities in the triples. We refer to this distant supervision process as the *alignment*
285 step (see Figure 2). The resulting text-triple pairs are subsequently refined in the *curation* step (see
286 Section 4.2) to retain only those pairs in which the textual content supports the associated TKGU
287 operations defined in Section 3.

288

289 4.2 QUALITY ASSURANCE AND CONTROL

290

291 During the *alignment* step of EMERGE creation pipeline (see Figure 2) we use multiple heuristics
292 to ensure the quality of the aligned textual passages with KG updates. For instance, we filter out
293 passages with a low proportion of English words and those containing wikitext special symbols used
294 for constructing elements such as tables and images. Furthermore, we discard updates in Wikidata
295 and Wikipedia that are quickly rolled back, as these often indicate incorrect or vandalized changes.
296 A complete list of preprocessing and cleaning steps can be found in Section L.4 in the appendix.

297 During the *Curation* step of the EMERGE pipeline (see Figure 2), we use
298 `Meta-Llama-3.1-405B` to validate that all TKGU operations can be derived from the
299 corresponding textual passage. **The full prompt design and illustrative examples are provided in**
300 **Appendix C.1.** This step flags KG updates not supported by the text, rather than removing them,
301 enabling future use of more powerful LLMs for additional verification and curation. Preserving
302 unsupported triples also allows evaluation of potential models that may rely less on text and more
303 on KG knowledge, particularly for EE-KG-Triples TKGU operations, where an entity may not
304 appear in the passage and updating the KG requires KG knowledge itself (e.g., all humans in the
305 KG link to the entity *human*). Appendix C.3 reports additional statistics on the fraction of triples
306 marked as unsupported.

307 Finally, during the *Curation* step, we manually annotate a random subset of 500 triple-text pairs (100
308 per TKGU operation type) to verify agreement with the LLM. We observe *Strong* to *Almost perfect*
309 agreement depending on the operation type, supporting the use of `Meta-Llama-3.1-405B` to
310 annotate the full dataset. Detailed annotation guidelines and agreement statistics are provided in
311 Appendices C.2.1 and C.2.2, respectively.

312

313 4.3 DATASET STATISTICS

314

315 EMERGE consists of 233K instances across seven yearly KG snapshots (2019–2025), with a total
316 of 1.45M TKGU update operations. Updates in each snapshot are evaluated over cumulative weekly
317 delta (Δ) intervals of up to 5 weeks. Both the KG size (i.e., number of entities and edges) and the
318 schema (i.e., number of relation types) evolve across snapshots. For instance, the 2019 KG snap-
319 shot contains 5.96M entities, 25.73M relations, and 5,646 relation types, while the 2025 snapshot
320 includes 6.93M entities, 37.54M relations, and 12,304 relation types. This dynamic setting enables
321 the evaluation of model robustness under evolving KG knowledge and schema changes, thereby *re-
322 flecting real-world KG evolution*. Additional tables and figures in Appendix D provide a detailed
323 overview of the size and distribution of TKGU operations in EMERGE. **Furthermore, Tables 7–12**
324 **in Appendix E present illustrative examples of each TKGU operation type introduced in Section 3.**

324 4.4 DATASET EXTENSION
325

326 EMERGE is an automatically constructed dataset, which we plan to extend using yearly snapshots
327 of Wikipedia and Wikidata, following the pipeline described in Section 4 and illustrated in Figure
328 2. These periodic extensions will enable the evaluation of architectures on their ability to extract
329 emerging real-world knowledge from text. This is particularly important for LLM-based architec-
330 tures, which are prone to hallucinating outdated information due to their internal parameters being
331 pre-trained on older textual sources (Wu et al., 2024a). To facilitate further development, we will
332 also provide code that allows users to extend the dataset themselves.

333
334

335 5 EXPERIMENTAL SETUP
336

337 We evaluate EMERGE using two state-of-the-art information extraction (IE) models that extract
338 structured knowledge as triples from text. These models are tested on a set constructed by subsam-
339 pling 5,000 instances from each snapshot (1,000 per delta), resulting in a total of 35,000 instances
340 and 201,369 TKGU operations. During subsampling, we retained up to 400 instances per delta con-
341 taining D-Triples TKGU operations. This ensures a sufficiently large number of D-Triples examples
342 for evaluation, even though they account for only 0.6% of all TKGU operations in the full dataset.
343 Conversely, in the test set, D-Triples constitute 3.3% (6,718 operations) of all TKGU operations.
344 **This low proportion of D-Triples does not affect metric stability, as each TKGU operation type is**
345 **evaluated independently rather than through aggregated performance across types (see Table 2).** A
346 detailed comparison of TKGU operation distributions is provided in Appendix D.2.

347
348

349 5.1 MODELS

350 To assess state-of-the-art performance on EMERGE, we evaluate two widely used IE architectures:
351 traditional extractive span-based models (Lee et al., 2017) and recent generative large language
352 models (LLMs) (Dagdelen et al., 2024; Xu et al., 2024; Zhang et al., 2025). For the span-based
353 setting, we use ReLiK (Orlando et al., 2024), and for the LLM-based setting, we adopt EDC (Zhang
354 & Soh, 2024). **Rather than comparing these models in terms of absolute performance, our goal is**
355 **to illustrate the complementary limitations of two mainstream IE paradigms when applied to text-
356 driven KG updating.** This setup highlights where each paradigm succeeds or fails across the different
357 TKGU operations defined in Section 3, particularly in their ability to handle emerging entities and
358 reason over existing KG structure. Below, we describe these architectures in more detail and explain
359 how we adapt them to each TKGU operation type.

360 **ReLiK.** ReLiK (Orlando et al., 2024) is a highly scalable architecture designed to minimize resource
361 usage while achieving state-of-the-art performance in both entity linking and relation extraction. In
362 our study, we evaluate two variants of ReLiK: closed information extraction ReLiK (ReLiK cIE)
363 and relation-extraction ReLiK (ReLiK RE). *ReLiK cIE* operates under the closed IE assumption
364 (Galárraga et al., 2014; Chaganty et al., 2017; Josifoski et al., 2023), predicting relations only be-
365 tween entities already present in the KG. Consequently, it can handle only those TKGU operations
366 involving known entities, namely, *X-Triples* and *E-Triples* as defined in Section 3. For each test
367 snapshot t , both models are provided with the corresponding KG snapshot. Specifically, ReLiK cIE
368 receives the dictionaries of entities (V_t) and relation types (R_t) present in t , while ReLiK RE is
369 given only the relation types (R_t), as it predicts relations without linking extracted entity mentions.
Further details on the ReLiK execution and configuration are provided in Appendix I.

370 **EDC.** The *extract, define, canonicalize* (EDC) framework, introduced by Zhang & Soh (2024),
371 is a state-of-the-art LLM-based approach. We adapt the original EDC prompt to additionally ex-
372 tract triples involving entities that are not explicitly mentioned in the input text but are poten-
373 tially present in a Wikidata KG snapshot. Furthermore, we extend this prompt even further, ask-
374 ing the model to identify potential triples to be deprecated from the KG. This way, we give the
375 model the ability to identify *EE-KG-Triples* and *D-Triples* operations based on the emerging ev-
376 idence in text (see Section 3). We term this adaptation **EDC+** in our experiments, and evaluate
377 it on *Mistral-7B-Instruct-v0.2* (*EDC+ Mistral-7b*) and *gemma-7b* (*EDC+ Gemma-7b*)
LLMs. Additional execution details as well as the used prompts are described in Appendix H.

378
 379 Table 2: *Recall* (measured using the completeness score) for IE models that do not link extracted
 380 triples to the KG, evaluated across KG snapshots on the TKGU operations defined in Section 3.

381 TKGU	382 Model	383 2019	384 2020	385 2021	386 2022	387 2023	388 2024	389 2025
390 X-Triples	391 EDC+ Mistral-7b	392 9.7	393 7.5	394 10.5	395 8.1	396 11.7	397 7.4	398 8.7
	399 EDC+ Gemma-7b	400 7.5	401 7.9	402 7.6	403 5.8	404 8.1	405 5.7	406 6.5
	407 ReLiK RE	408 25.3	409 24.5	410 24.1	411 20.1	412 22.1	413 19.2	414 20.3
415 E-Triples	416 EDC+ Mistral-7b	417 18.8	418 17.6	419 16.3	420 17.1	421 18.6	422 19.4	423 19.3
	424 EDC+ Gemma-7b	425 16.4	426 14.4	427 13.0	428 13.5	429 15.7	430 14.5	431 14.6
	432 ReLiK RE	433 23.3	434 20.3	435 23.1	436 15.9	437 17.0	438 15.0	439 16.4
440 EE-Triples	441 EDC+ Mistral-7b	442 21.3	443 16.7	444 10.0	445 15.7	446 18.4	447 13.2	448 15.6
	449 EDC+ Gemma-7b	450 18.4	451 13.5	452 9.1	453 14.8	454 17.2	455 13.0	456 13.2
	457 ReLiK RE	458 25.4	459 18.7	460 12.4	461 23.7	462 22.4	463 15.6	464 16.2
465 EE-KG-Triples	466 EDC+ Mistral-7b	467 25.6	468 19.9	469 7.1	470 23.0	471 21.6	472 16.9	473 18.3
	474 EDC+ Gemma-7b	475 11.3	476 8.0	477 2.0	478 8.2	479 8.9	480 7.3	481 6.5
	482 ReLiK RE	483 3.2	484 4.6	485 2.7	486 3.8	487 4.1	488 4.0	489 4.4
490 D-Triples	491 EDC+ Mistral-7b	492 7.1	493 9.8	494 7.7	495 7.7	496 12.6	497 4.0	498 8.7
	499 EDC+ Gemma-7b	500 5.5	501 10.6	502 8.4	503 10.4	504 10.8	505 5.8	506 6.7

396
 397 Table 3: *Recall* for the closed IE model ReLiK cIE (i.e., the extracted triples are linked to the KG)
 398 across KG snapshots, evaluated using the TKGU operations defined in Section 3.

400 TKGU Operations	401 Model	402 2019	403 2020	404 2021	405 2022	406 2023	407 2024	408 2025
409 X-Triples	410 ReLiK cIE	411 18.1	412 16.5	413 15.7	414 14.4	415 15.7	416 12.9	417 14.9
418 E-Triples	419 ReLiK cIE	420 14.9	421 16.8	422 14.0	423 13.4	424 15.2	425 12.5	426 14.7

427 5.2 METRICS AND EVALUATION

428 In order to evaluate the extraction and deprecation of triples based on emerging knowledge in text,
 429 we use recall as the primary metric (see Appendix B) to evaluate performance. We do not report
 430 precision or F1 scores in our main results Tables 2–3, as these metrics can be misleading under the
 431 open-world assumption (Razniewski et al., 2024). Under this assumption, the model may generate
 432 correct triple predictions that are incorrectly classified as false positives due to the inherently incom-
 433 plete nature of KGs, which do not necessarily capture the full set of valid triples. We additionally
 434 provide precision and F1 scores in Tables 13–16 in Appendix F.1.

435 For models that do not link extracted triples to KG, as indicated in the column *KG Link* in Table 1
 436 (ReLiK RE and EDC+), we evaluate recall with the *completeness score* (Jiang et al., 2024). This
 437 metric counts a ground-truth triple as correct if its cosine similarity with a predicted triple is above
 438 a set threshold (see Appendix B.2). This evaluation strategy is necessary because these models are
 439 not grounded in the entities present in the KG. This limitation underscores a broader research gap:
 440 existing IE methods operate largely independently of KG structure, making true text-driven KG
 441 updating challenging. Developing IE models that jointly exploit textual evidence and KG state to
 442 generate KG-grounded TKGU operations represents a promising direction for future work (see also
 443 Section K).

444 6 EXPERIMENTS AND ANALYSIS

445 Table 2 reports the performance of the ReLiK RE and EDC+ models across all TKGU operations.
 446 Table 3 shows the results for the ReLiK cIE model in the closed IE setting, which is restricted
 447 to TKGU operations involving existing entities and relations in the KG, namely *X-Triples* and *E-
 448 Triples*. The following paragraphs address key research questions and aim to lay the groundwork for
 449 future studies leveraging the TKGU operations introduced in this work.

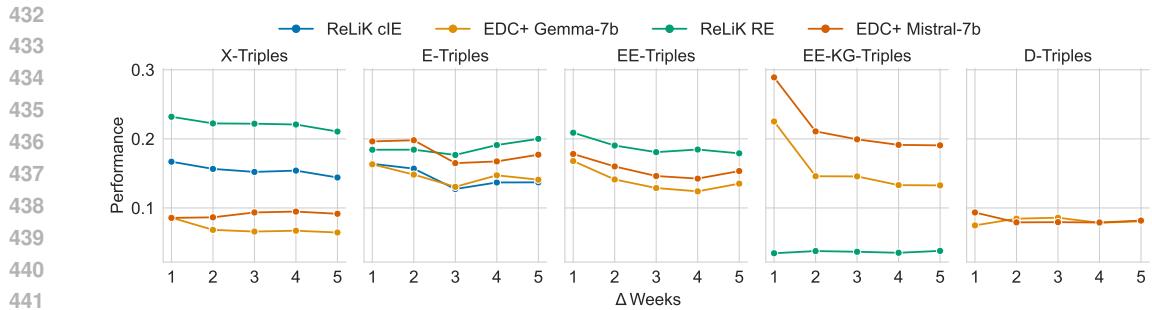


Figure 3: Performance of the models across temporal KG knowledge deltas. Some models show drops for certain TKGU operation types, for instance, EDC+Gemma-7b and EDC+Mistral-7b decline by over 5 percentage points between the first and second week deltas for EE-KG-Triples TKGU type.

What is the general performance? Overall, performance is low for both the recall metric reported in Table 3 and the completeness metric in Table 2. However, a closer inspection of the model predictions (see Appendix G) reveals that, in many cases, the extracted triples are semantically correct but do not align with the specific ground truth triples involved in the annotated TKGU operations. We hypothesize that this discrepancy arises because the models lack access to the KG content and structure, which prevents them from determining the nature of the knowledge being added and the types of relations involved. Access to KG-level statistics, such as the distribution of relation types, could provide valuable context and help improve model performance. This also points to a promising direction for future research: developing IE models that can identify emerging knowledge from unstructured text while leveraging the internal structure and temporal dynamics of the KGs.

How do ReLiK and EDC+ differ in handling TKGU operations? We selected the LLM-driven generative EDC+ model and the traditional, lightweight extractive span-based ReLiK model to compare how two fundamentally different and widely used architectures perform on TKGU operations. From Table 2, we observe that ReLiK RE significantly outperforms EDC+ on *X-Triples*. We hypothesize that this gap arises because ReLiK cIE and RE are explicitly trained to extract Wikidata triples from Wikipedia text, allowing the models to better capture relation structures and their distribution in the EMERGE corpora. In contrast, EDC+ relies only on a few in-context examples provided in the prompt, which appears insufficient to capture the diversity and complexity of relation types present in the dataset.

For TKGU operations that add previously non-existing triples to the KG, EDC+ performs comparably to ReLiK on *E-Triples* and *EE-Triples*. Furthermore, EDC+ significantly outperforms ReLiK RE on the *EE-KG-Triples* operation, which involves linking emerging entities mentioned in the passage to existing KG entities that are not explicitly referenced in the same passage. This result is expected, as ReLiK RE is designed to extract only entities explicitly mentioned in the text, as is also the case of other existing state-of-the-art IE models (see *EE-KG-Triples* column in Table 1). Its low performance on *EE-KG-Triples* is largely due to its reliance on explicit mentions: it extracts valid triples involving both emerging and existing entities that are present in the text but are not annotated in EMERGE, which includes only entity mentions explicitly annotated via Wikipedia hyperlinks.

When evaluated on *D-Triples*, EDC+ demonstrates relatively low performance, largely due to its lack of access to the knowledge graph. Without this information, the model cannot reliably identify triples that are already present and should be deprecated (see Appendix G for an example). In contrast, ReLiK is not trained to explicitly identify triples to be removed from KG and therefore is unable to extract *D-Triples* TKGU operation. This limitation also applies to other state-of-the-art IE models (see *D-Triples* column in Table 1).

What is the performance across different snapshots? Although results on earlier snapshots appear slightly higher than those from later years across different models and TKGU operations, there is no clear overall trend. We hypothesize that these performance differences are driven less by the

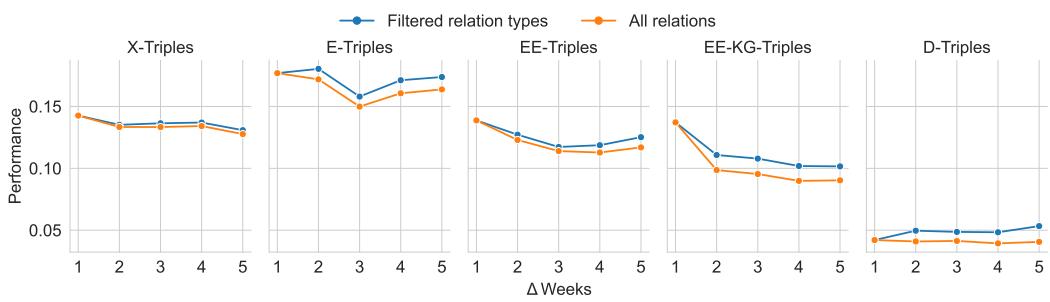


Figure 5: Performance of TKGU operations on relation types from the first KG delta week (*Filtered relations*) versus the full dataset including all relation types (*All relations*). The increased performance on *Filtered relations* shows that newly introduced relation types in later deltas are harder to predict, leading to larger performance drops.

novelty of the knowledge itself and more by the type of emerging knowledge dominant in each snapshot, an aspect we plan to investigate in future work.

What is the performance on increasing temporal KG deltas? In Figure 3, we plot model performance across increasing weekly KG deltas. Although not consistent across all TKGU operations and models, we generally observe a performance drop as deltas grow. **We hypothesize that this decline stems from an increased number of relation types involved in the TKGU operations at higher deltas (see Figure 4).** To test this hypothesis, we evaluate TKGU operations from the knowledge delta of week 2 onward while restricting relation types to those already present in week 1. Figure 5 shows the average performance difference across the evaluated models as the delta interval grows. Here, *Filtered relation types* denote performance restricted to relation types seen in week 1, while *All relations* corresponds to performance on the full set of relation types at each update. The reduced performance drop in the filtered setting supports our claim. In future work, we plan to further investigate this phenomenon and develop more robust models for continual knowledge updates under ever-increasing temporal deltas.

7 CONCLUSION

In this work, we introduced EMERGE, the first dataset to cover all text-driven knowledge graph updating (TKGU) operations required to keep KGs aligned with emerging knowledge from textual sources. Evaluation of two state-of-the-art models on a dataset subset revealed a gap in current information extraction models to extract new information from text while accounting for existing KG content and structure. This suggests that future work should focus on designing methods capable of interacting with both emerging knowledge in text and the evolving content and structure of KGs. Additional limitations of our work, along with potential directions for future research, are discussed in Appendix K.

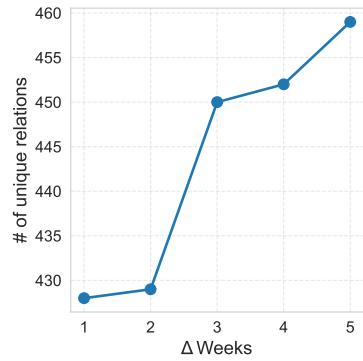


Figure 4: Evolution of the number of relation types with increasing weekly KG deltas.

540 8 REPRODUCIBILITY STATEMENT
541

542 The code for dataset creation and reproducing the experimental results will be released in a public
543 GitHub repository. The repository will also provide functionality for extending EMERGE with new
544 KG snapshots, enabling incorporation of novel emerging knowledge (see Section 4.4). Moreover,
545 the LLMs used for dataset annotation (Section 4.2) and within the EDC+ model (Section 5.1) are
546 publicly accessible, enabling straightforward replication of dataset construction and experiments.
547

548 9 ETHICS STATEMENT
549

550 We confirm that we have read and adhered to the ICLR Code of Ethics throughout this work. Our
551 study does not involve human subjects, personally identifiable information, or sensitive data (refer
552 to Appendix L.2 for further details), and no ethical approval (e.g., IRB) was required. The datasets
553 used are publicly available and comply with licensing and privacy requirements. We are not aware
554 of any potential harms, security risks, or fairness concerns arising from the methods or applications
555 of our research. There are no conflicts of interest, sponsorship influences, or legal compliance issues
556 to disclose.
557

558 REFERENCES
559

560 Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. Knowledge graph based syn-
561 thetic corpus generation for knowledge-enhanced language model pre-training. In *Proceed-
562 ings of the 2021 Conference of the North American Chapter of the Association for Com-
563 putational Linguistics: Human Language Technologies (NAACL-HLT 2021)*, pp. 3554–3565,
564 2021. doi: 10.18653/v1/2021.naacl-main.278. URL <https://aclanthology.org/2021.naacl-main.278>.

565 Garima Agrawal, Tharindu Kumarage, Zeyad Alghamdi, and Huan Liu. Can knowledge graphs
566 reduce hallucinations in llms?: A survey. In *Proceedings of the 2024 Conference of the North
567 American Chapter of the Association for Computational Linguistics: Human Language Tech-
568 nologies (Volume 1: Long Papers)*, pp. 3947–3960, 2024.

569 Jacqueline Aguilar, Charley Beller, Paul McNamee, Benjamin Van Durme, Stephanie Strassel, Zhiyi
570 Song, and Joe Ellis. A comparison of the events and relations across ace, ere, tac-kbp, and
571 framenet annotation standards. In *Proceedings of the 2nd Workshop on EVENTS: Defini-
572 tion, De-
573 tection, Coreference, and Representation*, pp. 45–53, 2014. URL <https://aclanthology.org/W14-2907.pdf>.

575 Mehwish Alam, Genet Asefa Gesese, and Pierre-Henri Paris. Neurosymbolic methods for dynamic
576 knowledge graphs. *arXiv preprint arXiv:2409.04572*, 2024.

577 Isabelle Augenstein, Mrinal Das, Sebastian Riedel, Lakshmi Vikraman, and Andrew McCallum.
578 Semeval 2017 task 10: Scienceie-extracting keyphrases and relations from scientific publica-
579 tions. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-
580 2017)*, pp. 546–555, 2017. doi: 10.18653/v1/S17-2091. URL <https://aclanthology.org/S17-2091>.

582 Zhen Bi, Jing Chen, Yinuo Jiang, Feiyu Xiong, Wei Guo, Huajun Chen, and Ningyu Zhang.
583 Codekgc: Code language model for generative knowledge graph construction. *ACM Transac-
584 tions on Asian and Low-Resource Language Information Processing*, 23(3):1–16, 2024. doi:
585 10.1145/3641850. URL <https://dl.acm.org/doi/full/10.1145/3641850>.

586 Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: a collab-
587 oratively created graph database for structuring human knowledge. In *Proceedings of the 2008
588 ACM SIGMOD international conference on Management of data*, pp. 1247–1250, 2008. doi:
589 10.1145/1376616.1376746. URL <https://doi.org/10.1145/1376616.1376746>.

591 Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko.
592 Translating embeddings for modeling multi-relational data. In *Advances in neural information
593 processing systems*, pp. 2787–2795, 2013. URL <https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html>.

594 Elizabeth Boschee, Jennifer Lautenschlager, Sean O'Brien, Steve Shellman, James Starz,
595 and Michael Ward. ICEWS coded event data. *Harvard Dataverse*, 12, 2015. URL
596 <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/28075>.
597

598 Pere-Lluís Huguet Cabot and Roberto Navigli. REBEL: Relation extraction by end-to-end lan-
599 guage generation. In *Findings of the Association for Computational Linguistics: EMNLP*
600 2021, pp. 2370–2381, 2021. doi: 10.18653/v1/2021.findings-emnlp.204. URL <https://aclanthology.org/2021.findings-emnlp.204/>.
601

602 Linyue Cai, Chaojia Yu, Yongqi Kang, Yu Fu, Heng Zhang, and Yong Zhao. Practices, opportuni-
603 ties and challenges in the fusion of knowledge graphs and large language models. *Frontiers in*
604 *Computer Science*, 7:1590632, 2025.
605

606 Arun Chaganty, Ashwin Paranjape, Percy Liang, and Christopher D Manning. Importance sampling
607 for unbiased on-demand evaluation of knowledge base population. In *Proceedings of the 2017*
608 *Conference on Empirical Methods in Natural Language Processing*, pp. 1038–1048, 2017. doi:
609 10.18653/v1/D17-1109. URL <https://aclanthology.org/D17-1109/>.
610

611 Hanzhu Chen, Xu Shen, Jie Wang, Zehao Wang, Qitan Lv, Junjie He, Rong Wu, Feng Wu, and
612 Jieping Ye. Knowledge graph finetuning enhances knowledge manipulation in large language
613 models. In *The Thirteenth International Conference on Learning Representations*, 2025.
614

615 Liyi Chen, Panrong Tong, Zhongming Jin, Ying Sun, Jieping Ye, and Hui Xiong. Plan-on-graph:
616 Self-correcting adaptive planning of large language model on knowledge graphs. *Advances in*
617 *Neural Information Processing Systems*, 37:37665–37691, 2024.
618

619 Nancy Chinchor and Elaine Marsh. Muc-7 information extraction task definition. In *Proceeding*
620 *of the 1998 Message Understanding Conference (MUC-7)*, pp. 359–367, 1998. URL <https://catalog.ldc.upenn.edu/docs/LDC2001T02/guidelines.IEtask42.ps>.
621

622 John Dagdelen, Alexander Dunn, Sanghoon Lee, Nicholas Walker, Andrew S Rosen, Gerbrand
623 Ceder, Kristin A Persson, and Anubhav Jain. Structured information extraction from sci-
624 entific text with large language models. *Nature Communications*, 15(1):1418, 2024. doi: <https://doi.org/10.1038/s41467-024-45563-x>. URL <https://www.nature.com/articles/s41467-024-45563-x>.
625

626 Shib Sankar Dasgupta, Swayambhu Nath Ray, and Partha Talukdar. Hyte: Hyperplane-based tem-
627 porally aware knowledge graph embedding. In *Proceedings of the 2018 conference on empirical*
628 *methods in natural language processing*, pp. 2001–2011, 2018. doi: 10.18653/v1/D18-1225.
629 URL <https://aclanthology.org/D18-1225/>.
630

631 Daniel Daza, Michael Cochez, and Paul Groth. Inductive entity representations from text via link
632 prediction. In *Proceedings of the Web Conference 2021*, pp. 798–808, 2021. doi: 10.1145/
633 3442381.3450141. URL <https://doi.org/10.1145/3442381.3450141>.
634

635 Leon Derczynski, Eric Nichols, Marieke van Erp, and Nut Limsopatham. Results of the WNUT2017
636 shared task on novel and emerging entity recognition. In *Proceedings of the 3rd Workshop on*
637 *Noisy User-generated Text*, pp. 140–147, 2017. doi: 10.18653/V1/W17-4418. URL <https://doi.org/10.18653/v1/w17-4418>.
638

639 Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d
640 knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*,
641 volume 32, 2018. doi: <https://doi.org/10.1609/aaai.v32i1.11573>. URL <https://ojs.aaai.org/index.php/AAAI/article/view/11573>.
642

643 Bhuwan Dhingra, Jeremy R Cole, Julian Martin Eisenschlos, Daniel Gillick, Jacob Eisenstein, and
644 William W Cohen. Time-aware language models as temporal knowledge bases. *Transactions*
645 *of the Association for Computational Linguistics*, 10:257–273, 2022. doi: 10.1162/tacl_a_00459.
646 URL https://doi.org/10.1162/tacl_a_00459.
647

648 Bayu Distiawan, Gerhard Weikum, Jianzhong Qi, and Rui Zhang. Neural relation extraction for
649 knowledge base enrichment. In *Proceedings of the 2019 Conference of the Association for*
650 *Computational Linguistics*, pp. 229–240, 2019. doi: 10.18653/v1/P19-1023. URL <https://aclanthology.org/P19-1023/>.

651

652 Zixuan Dong, Baoyun Peng, Yufei Wang, Jia Fu, Xiaodong Wang, Xin Zhou, Yongxue Shan,
653 Kangchen Zhu, and Weigu Chen. Effiqa: Efficient question-answering with strategic multi-
654 model collaboration on knowledge graphs. In *Proceedings of the 31st International Conference*
655 *on Computational Linguistics*, pp. 7180–7194, 2025.

656

657 Karel D’Oosterlinck, François Remy, Johannes Deleu, Thomas Demeester, Chris Develder, Klim
658 Zaporojets, Arya Ghodsi, S Ellershaw, J Collins, and C Potts. BioDEX: Large-scale biomedical
659 adverse drug event extraction for real-world pharmacovigilance. In *Findings of the ACL,*
660 *will be held at EMNLP 2023, The 2023 Conference on Empirical Methods in Natural Lan-*
661 *guage Processing*, 2023. doi: 10.18653/V1/2023.FINDINGS-EMNLP.896. URL <https://doi.org/10.18653/v1/2023.findings-emnlp.896>.

662

663 Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt,
664 and Jonathan Larson. From local to global: A graph rag approach to query-focused summariza-
665 tion. *arXiv preprint arXiv:2404.16130*, 2024.

666

667 Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Elena Sim-
668 perl, and Frederique Laforest. T-REx: A large scale alignment of natural language with knowledge
669 base triples. 2019. URL <http://www.lrec-conf.org/proceedings/lrec2018/summaries/632.html>.

670

671 Evgeniy Gabrilovich, Michael Ringgaard, and Amarnag Subramanya. FACC1: Freebase annotation
672 of clueweb corpora, 2013. URL <https://lemurproject.org/clueweb12/FACC1/>.

673

674 Luis Galárraga, Jeremy Heitz, Kevin Murphy, and Fabian M Suchanek. Canonicalizing open knowl-
675 edge bases. In *Proceedings of the 23rd ACM international conference on conference on informa-*
676 *tion and knowledge management*, pp. 1679–1688, 2014. doi: 10.1145/2661829.2662073. URL
677 <https://doi.org/10.1145/2661829.2662073>.

678

679 Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. Fewrel 2.0:
680 Towards more challenging few-shot relation classification. In *Proceedings of the 2019 Confer-*
681 *ence on Empirical Methods in Natural Language Processing and the 9th International Joint*
682 *Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 6250–6255, 2019. doi:
683 10.18653/v1/D19-1649. URL <https://aclanthology.org/D19-1649/>.

684

685 Alberto Garcia-Duran, Sebastijan Dumančić, and Mathias Niepert. Learning sequence encoders
686 for temporal knowledge graph completion. In *Proceedings of the 2018 Conference on Empirical*
687 *Methods in Natural Language Processing*, pp. 4816–4821, 2018.

688

689 Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. The webnlg
690 challenge: Generating text from rdf data. In *Proceedings of the 10th International Conference on*
691 *Natural Language Generation*, pp. 124–133, 2017. doi: 10.18653/v1/W17-3518. URL <https://aclanthology.org/W17-3518/>.

692

693 Julia Gastinger, Shenyang Huang, Michael Galkin, Erfan Loghmani, Ali Parviz, Farimah Poursafaei,
694 Jacob Danovitch, Emanuele Rossi, Ioannis Koutis, Heiner Stuckenschmidt, et al. Tgb 2.0: A
695 benchmark for learning on temporal knowledge graphs and heterogeneous graphs. *Advances in*
696 *neural information processing systems*, 37:140199–140229, 2024.

697

698 Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach,
699 Hal Daumé III, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64
700 (12):86–92, 2021. doi: 10.1145/3458723. URL <https://doi.org/10.1145/3458723>.

701

Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. Grammar-constrained decoding
702 for structured nlp tasks without finetuning. In *Proceedings of the 2023 Conference on Empiri-*
703 *cal Methods in Natural Language Processing*, pp. 10932–10952, 2023. doi: 10.18653/v1/2023.
704 emnlp-main.674. URL <https://aclanthology.org/2023.emnlp-main.674/>.

702 Genet Asefa Gesese, Mehwish Alam, and Harald Sack. Literallywikidata-a benchmark for knowl-
703 edge graph completion using literals. In *International Semantic Web Conference*, pp. 511–527.
704 Springer, 2021.

705 Jeremy Getman, Joe Ellis, Zhiyi Song, Jennifer Tracey, and Stephanie M Strassel.
706 Overview of linguistic resources for the tac kbp 2017 evaluations: Methodologies
707 and results. In *Proceedings of the 2017 Text Analysis Conference (TAC 2017)*, 2017.
708 URL https://tac.nist.gov/publications/2017/additional.papers/TAC2017.KBP_resources_overview.proceedings.pdf.

709

710 Harsha Gurulingappa, Abdul Mateen Rajput, Angus Roberts, Juliane Fluck, Martin Hofmann-
711 Apitius, and Luca Toldo. Development of a benchmark corpus to support the automatic extraction
712 of drug-related adverse effects from medical case reports. *Journal of biomedical informatics*, 45
713 (5):885–892, 2012. doi: 10.1016/J.JBI.2012.04.008. URL <https://doi.org/10.1016/j.jbi.2012.04.008>.

714

715 Haoyu Han, Yu Wang, Harry Shomer, Kai Guo, Jiayuan Ding, Yongjia Lei, Mahantesh Halap-
716 panavar, Ryan A Rossi, Subhabrata Mukherjee, Xianfeng Tang, et al. Retrieval-augmented gen-
717 eration with graphs (graphrag). *arXiv preprint arXiv:2501.00309*, 2024.

718

719 Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. FewRel:
720 A large-scale supervised few-shot relation classification dataset with state-of-the-art evalua-
721 tion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Pro-
722 cessing (EMNLP 2018)*, pp. 4803–4809, 2018. doi: 10.18653/v1/D18-1514. URL <https://aclanthology.org/D18-1514>.

723

724

725 Xu Han, Tianyu Gao, Yuan Yao, Deming Ye, Zhiyuan Liu, and Maosong Sun. OpenNRE: An
726 open and extensible toolkit for neural relation extraction. In *Proceedings of the 2019 Conference
727 on Empirical Methods in Natural Language Processing and the 9th International Joint Confer-
728 ence on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pp. 169–174,
729 2019. doi: 10.18653/V1/D19-3029. URL <https://doi.org/10.18653/v1/D19-3029>.

730

731 Xu Han, Tianyu Gao, Yankai Lin, Hao Peng, Yaoliang Yang, Chaojun Xiao, Zhiyuan Liu, Peng
732 Li, Jie Zhou, and Maosong Sun. More data, more relations, more context and more open-
733 ness: A review and outlook for relation extraction. In *Proceedings of the 1st Conference of the
734 Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International
735 Joint Conference on Natural Language Processing*, pp. 745–758, 2020. doi: 10.18653/V1/2020.
AACL-MAIN.75. URL <https://doi.org/10.18653/v1/2020.aacl-main.75>.

736

737 Yuanzhen Hao and Desheng Wu. Fact verification on knowledge graph via programmatic graph
738 reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp.
5480–5495, 2025.

739

740 Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Se-
741 bastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. Semeval-2010
742 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Pro-
743 ceedings of the 5th International Workshop on Semantic Evaluation*, pp. 33–38, 2010. URL
744 <https://aclanthology.org/S10-1006.pdf>.

745

746 María Herrero-Zazo, Isabel Segura-Bedmar, Paloma Martínez, and Thierry Declerck. The DDI cor-
747 pus: An annotated corpus with pharmacological substances and drug–drug interactions. *Jour-
748 nal of biomedical informatics*, 46(5):914–920, 2013. doi: 10.1016/J.JBI.2013.07.011. URL
<https://doi.org/10.1016/j.jbi.2013.07.011>.

749

750 Daniel Hewlett, Alexandre Lacoste, Llion Jones, Illia Polosukhin, Andrew Fandrianto, Jay Han,
751 Matthew Kelcey, and David Berthelot. Wikireading: A novel large-scale language understanding
752 task over wikipedia. In *Proceedings of the 54th Annual Meeting of the Association for Compu-
753 tational Linguistics (Volume 1: Long Papers)*, pp. 1535–1545, 2016. doi: 10.18653/V1/P16-1145.
754 URL <https://doi.org/10.18653/v1/p16-1145>.

755 Marvin Hofer, Daniel Obraczka, Alieh Saeedi, Hanna Köpcke, and Erhard Rahm. Construction of
knowledge graphs: Current state and challenges. *Information*, 15(8):509, 2024.

756 Yutai Hou, Yingce Xia, Lijun Wu, Shufang Xie, Yang Fan, Jinhua Zhu, Tao Qin, and Tie-Yan Liu.
757 Discovering drug–target interaction knowledge from biomedical literature. *Bioinformatics*, 38
758 (22):5100–5107, 2022. doi: 10.1093/BIOINFORMATICS/BTAC648. URL <https://doi.org/10.1093/bioinformatics/btac648>.

760 Yujia Hu, Tuan-Phong Nguyen, Shrestha Ghosh, Moritz Müller, and Simon Razniewski. Gp-
761 tkb v1. 5: A massive knowledge base for exploring factual llm knowledge. *arXiv preprint*
762 *arXiv:2507.05740*, 2025.

763 Shenyang Huang, Farimah Poursafaei, Jacob Danovitch, Matthias Fey, Weihua Hu, Emanuele Rossi,
764 Jure Leskovec, Michael Bronstein, Guillaume Rabusseau, and Reihaneh Rabbany. Temporal
765 graph benchmark for machine learning on temporal graphs. *Advances in Neural Information
766 Processing Systems*, 36, 2024.

767 Joel Jang, Seonghyeon Ye, Changho Lee, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun
768 Kim, and Minjoon Seo. TemporalWiki: A lifelong benchmark for training and evaluating ever-
769 evolving language models. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Pro-
770 ceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp.
771 6237–6250, Abu Dhabi, United Arab Emirates, December 2022a. Association for Computational
772 Linguistics. doi: 10.18653/v1/2022.emnlp-main.418. URL <https://aclanthology.org/2022.emnlp-main.418/>.

773 Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, Stan-
774 ley Jungkyu Choi, and Minjoon Seo. Towards continual knowledge learning of language models.
775 In *ICLR*, 2022b. URL <https://openreview.net/forum?id=vfsRB5MImo9>.

776 Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. Overview of the TAC
777 2010 knowledge base population track. In *Proceedings of the 2010 Text Analysis Confer-
778 ence (TAC 2010)*, pp. 1–25, 2010. URL <https://blender.cs.illinois.edu/paper/kbp2010overview.pdf>.

779 Pengcheng Jiang, Jiacheng Lin, Zifeng Wang, Jimeng Sun, and Jiawei Han. Genres: Rethinking
780 evaluation for generative relation extraction in the era of large language models. In *Proceed-
781 ings of the 2024 Conference of the North American Chapter of the Association for Compu-
782 tational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 2820–2837,
783 2024. doi: 10.18653/v1/2024.naacl-long.155. URL <https://aclanthology.org/2024.naacl-long.155/>.

784 Pengcheng Jiang, Cao Xiao, Minhao Jiang, Parminder Bhatia, Taha Kass-Hout, Jimeng Sun, and
785 Jiawei Han. Reasoning-enhanced healthcare predictions with knowledge graph community re-
786 trieval. In *The Thirteenth International Conference on Learning Representations*, 2025.

787 Woojeong Jin, Meng Qu, Xisen Jin, and Xiang Ren. Recurrent event network: Autoregressive
788 structure inference over temporal knowledge graphs. In *Proceedings of the 2020 Conference on
789 Empirical Methods in Natural Language Processing (EMNLP)*, pp. 6669–6683, 2020.

790 Xiongnan Jin, Zhilin Wang, Manni Duan, Yan Shao, Xingyun Hong, Yongheng Wang, and
791 Byungkook Oh. A survey on knowledge graph evolution: proliferation, dynamic embedding,
792 and versioning. *International Journal of Web and Grid Services*, 21(1):88–111, 2025.

793 Martin Josifoski, Nicola De Cao, Maxime Peyrard, Fabio Petroni, and Robert West. GenIE: Genera-
794 tive information extraction. In Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir
795 Meza Ruiz (eds.), *Proceedings of the 2022 Conference of the North American Chapter of the As-
796 sociation for Computational Linguistics: Human Language Technologies*, pp. 4626–4643, Seat-
797 le, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.
798 naacl-main.342. URL <https://aclanthology.org/2022.naacl-main.342/>.

799 Martin Josifoski, Marija Sakota, Maxime Peyrard, and Robert West. Exploiting asymmetry for
800 synthetic training data generation: SynthIE and the case of information extraction. In Houda
801 Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical
802 Methods in Natural Language Processing*, pp. 1555–1574, Singapore, December 2023. Asso-
803 ciation for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.96. URL <https://aclanthology.org/2023.emnlp-main.96/>.

810 Jungo Kasai, Keisuke Sakaguchi, Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir
811 Radev, Noah A Smith, Yejin Choi, Kentaro Inui, et al. Realtime qa: What's the an-
812 swer right now? *Advances in Neural Information Processing Systems*, 36, 2024. URL
813 https://proceedings.neurips.cc/paper_files/paper/2023/file/9941624ef7f867a502732b5154d30cb7-Paper-Datasets_and_Benchmarks.pdf.

814

815

816 Nora Kassner, Philipp Dufter, and Hinrich Schütze. Multilingual LAMA: Investigating knowledge
817 in multilingual pretrained language models. In *Proceedings of the 2021 Conference of the Eu-
818 ropean Chapter of the Association for Computational Linguistics (EACL 2021)*, pp. 3250–3258,
819 2021. doi: 10.18653/V1/2021.EACL-MAIN.284. URL <https://doi.org/10.18653/v1/2021.eacl-main.284>.

820

821 Simerjot Kaur, Charese Smiley, Akshat Gupta, Joy Sain, Dongsheng Wang, Suchetha Siddagan-
822 gappa, Toyin Aguda, and Sameena Shah. REFinD: Relation extraction financial dataset. In
823 *Proceedings of the 46th international ACM SIGIR conference on research and development in
824 information retrieval*, pp. 3054–3063, 2023. doi: 10.1145/3539618.3591911. URL <https://doi.org/10.1145/3539618.3591911>.

825

826

827 Jiho Kim, Sungjin Park, Yeonsu Kwon, Yohan Jo, James Thorne, and Edward Choi. Factkg: Fact
828 verification via reasoning on knowledge graphs. *arXiv preprint arXiv:2305.06590*, 2023.

829

830 Jinyoung Kim, Dayoon Ko, and Gunhee Kim. DynamicER: Resolving emerging mentions to dy-
831 namic entities for RAG. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Pro-
832 ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp.
833 13752–13770, Miami, Florida, USA, November 2024a. Association for Computational Linguis-
834 tics. doi: 10.18653/v1/2024.emnlp-main.762. URL <https://aclanthology.org/2024.emnlp-main.762/>.

835

836 Yujin Kim, Jaehong Yoon, Seonghyeon Ye, Sangmin Bae, Namgyu Ho, Sung Ju Hwang, and Se-
837 Young Yun. Carpe diem: On the evaluation of world knowledge in lifelong language models.
838 In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Confer-
839 ence of the North American Chapter of the Association for Computational Linguistics: Human
840 Language Technologies (Volume 1: Long Papers)*, pp. 5401–5415, Mexico City, Mexico, June
841 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.302. URL
842 <https://aclanthology.org/2024.naacl-long.302/>.

843

844 Dayoon Ko, Jinyoung Kim, Hahyeon Choi, and Gunhee Kim. GrowOVER: How can LLMs
845 adapt to growing real-world knowledge? In Lun-Wei Ku, Andre Martins, and Vivek Sriku-
846 mar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Lin-
847 guistics (Volume 1: Long Papers)*, pp. 3282–3308, Bangkok, Thailand, August 2024. Asso-
848 ciation for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.181. URL <https://aclanthology.org/2024.acl-long.181/>.

849

850 Ruben Kruiper, Julian Vincent, Jessica Chen-Burger, Marc Desmulliez, and Ioannis Konstas.
851 In layman's terms: Semi-open relation extraction from scientific texts. In *Proceedings of
852 the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1489–1500,
853 2020. doi: 10.18653/V1/2020.ACL-MAIN.137. URL <https://doi.org/10.18653/v1/2020.acl-main.137>.

854

855 Timothée Lacroix, Guillaume Obozinski, and Nicolas Usunier. Tensor decompositions for temporal
856 knowledge base completion. In *International Conference on Learning Representations (ICLR)*,
857 2020. URL <https://openreview.net/forum?id=rke2P1BFwS>.

858

859 Ernests Lavrinovics, Russa Biswas, Johannes Bjerva, and Katja Hose. Knowledge graphs, large
860 language models, and hallucinations: An nlp perspective. *Journal of Web Semantics*, 85:100844,
861 2025.

862

863 Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. End-to-end neural coreference resolu-
864 tion. In Martha Palmer, Rebecca Hwa, and Sebastian Riedel (eds.), *Proceedings of the 2017 Con-
865 ference on Empirical Methods in Natural Language Processing*, pp. 188–197, Copenhagen, Den-
866 mark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1018.
867 URL <https://aclanthology.org/D17-1018/>.

864 Kalev Leetaru and Philip A Schrod. Gdelt: Global data on events, location, and tone, 1979–2012.
865 In *ISA annual convention*, number 4, pp. 1–49, 2013.

866

867 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Na-
868 man Goyal, Heinrich Kütterer, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al.
869 Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of*
870 *the 2020 Advances in Neural Information Processing Systems (NeurIPS 2020)*, pp.
871 9459–9474, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>.

872

873 Belinda Z Li, Emmy Liu, Alexis Ross, Abbas Zeitoun, Graham Neubig, and Jacob Andreas. Lan-
874 guage modeling with editable external knowledge. In *Findings of the Association for Compu-
875 tational Linguistics: NAACL 2025*, pp. 3070–3090, 2025a. doi: 10.18653/v1/2025.findings-naacl.
876 168. URL <https://aclanthology.org/2025.findings-naacl.168/>.

877 Jiao Li, Yueping Sun, Robin J Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Al-
878 lan Peter Davis, Carolyn J Mattingly, Thomas C Wiegers, and Zhiyong Lu. Biocreative v cdr
879 task corpus: a resource for chemical disease relation extraction. *Database*, 2016, 2016. doi:
880 10.1093/database/baw068. URL <https://doi.org/10.1093/database/baw068>.

881 Mengran Li, Pengyu Zhang, Wenbin Xing, Yijia Zheng, Klim Zaporojets, Junzhou Chen, Ronghui
882 Zhang, Yong Zhang, Siyuan Gong, Jia Hu, et al. A survey of large language models for data
883 challenges in graphs. *Expert Systems with Applications*, pp. 129643, 2025b.

884

885 Mufei Li, Siqi Miao, and Pan Li. Simple is effective: The roles of graphs and large language mod-
886 els in knowledge-graph-based retrieval-augmented generation. In *The Thirteenth International
887 Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*, 2025c. URL
888 <https://openreview.net/forum?id=JvkuZZ0407>.

889 Ke Liang, Lingyuan Meng, Meng Liu, Yue Liu, Wenxuan Tu, Siwei Wang, Sihang Zhou, Xinwang
890 Liu, Fuchun Sun, and Kunlun He. A survey of knowledge graph reasoning on graph types: Static,
891 dynamic, and multi-modal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46
892 (12):9456–9478, 2024.

893 Adam Liska, Tomas Kociský, Elena Gribovskaya, Tayfun Terzi, Eren Sezener, Devang Agrawal,
894 D'Autume Cyprien De Masson, Tim Scholtes, Manzil Zaheer, Susannah Young, et al. Stream-
895 ingQA: A benchmark for adaptation to new knowledge over time in question answering mod-
896 els. In *International Conference on Machine Learning*, pp. 13604–13622. PMLR, 2022. URL
897 <https://proceedings.mlr.press/v162/liska22a.html>.

898 Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. Multi-task identification of entities,
899 relations, and coreference for scientific knowledge graph construction. In *Proceedings of the 2018
900 Conference on Empirical Methods in Natural Language Processing*, pp. 3219–3232, 2018. doi:
901 10.18653/v1/D18-1360. URL <https://aclanthology.org/D18-1360/>.

902 Ling Luo, Po-Ting Lai, Chih-Hsuan Wei, Cecilia N Arighi, and Zhiyong Lu. BioRED: a rich
903 biomedical relation extraction dataset. *Briefings in Bioinformatics*, 23(5):bbac282, 2022. doi:
904 10.1093/bib/bbac282. URL <https://doi.org/10.1093/bib/bbac282>.

905 Linhao Luo, Zicheng Zhao, Gholamreza Haffari, Yuan-Fang Li, Chen Gong, and Shirui Pan. Graph-
906 constrained reasoning: Faithful reasoning on knowledge graphs with large language models. In
907 *Forty-second International Conference on Machine Learning*, 2025.

908

909 Xin Lv, Yankai Lin, Yixin Cao, Lei Hou, Juanzi Li, Zhiyuan Liu, Peng Li, and Jie Zhou. Do
910 pre-trained models benefit knowledge graph completion? a reliable evaluation and a reasonable
911 approach. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Findings of*
912 *the Association for Computational Linguistics: ACL 2022*, pp. 3570–3581, Dublin, Ireland, May
913 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-acl.282. URL
914 <https://aclanthology.org/2022.findings-acl.282/>.

915 Jie Ma, Ning Qu, Zhitao Gao, Rui Xing, Jun Liu, Hongbin Pei, Jiang Xie, Linyun Song, Pinghui
916 Wang, Jing Tao, et al. Deliberation on priors: Trustworthy reasoning of large language models
917 on knowledge graphs. In *The Thirty-ninth Annual Conference on Neural Information Processing
Systems*, 2025.

918 Farzaneh Mahdisoltani, Joanna Biega, and Fabian Suchanek. Yago3: A knowledge base from mul-
919 tilingual wikipedias. In *7th biennial conference on innovative data systems research*. CIDR Con-
920 ference, 2014. URL <https://imt.hal.science/hal-01699874/>.

921

922 Katerina Margatina, Shuai Wang, Yogarshi Vyas, Neha Anna John, Yassine Benajiba, and Miguel
923 Ballesteros. Dynamic benchmarking of masked language models on temporal concept drift with
924 multiple views. In Andreas Vlachos and Isabelle Augenstein (eds.), *Proceedings of the 17th*
925 *Conference of the European Chapter of the Association for Computational Linguistics*, pp. 2881–
926 2898, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/
927 v1/2023.eacl-main.211. URL <https://aclanthology.org/2023.eacl-main.211/>.

928

929 Filipe Mesquita, Matteo Cannaviccio, Jordan Schmidek, Paramita Mirza, and Denilson Barbosa.
930 KnowledgeNet: A benchmark dataset for knowledge base population. In Kentaro Inui, Jing
931 Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empir-
932 ical Methods in Natural Language Processing and the 9th International Joint Conference on
933 Natural Language Processing (EMNLP-IJCNLP)*, pp. 749–758, Hong Kong, China, November
934 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1069. URL
<https://aclanthology.org/D19-1069/>.

935

936 Nandana Mihindukulasooriya, Sanju Tiwari, Carlos F Enguix, and Kusum Lata. Text2KG Bench: A
937 benchmark for ontology-driven knowledge graph generation from text. In *International Semantic
938 Web Conference*, pp. 247–265. Springer, 2023. doi: 10.1007/978-3-031-47243-5_14. URL
https://doi.org/10.1007/978-3-031-47243-5_14.

939

940 George A Miller. WordNet: a lexical database for english. *Communications of the ACM*, 38(11):
941 39–41, 1995. doi: 10.1145/219717.219748. URL <https://doi.org/10.1145/219717.219748>.

942

943 Antonio Miranda, Farrokh Mehryary, Jouni Luoma, Sampo Pyysalo, Alfonso Valencia, and Mar-
944 tin Krallinger. Overview of drugprot biocreative vii track: quality evaluation and large scale
945 text mining of drug-gene/protein relations. In *Proceedings of the seventh BioCreative challenge*
946 *evaluation workshop*, pp. 11–21, 2021.

947

948 Yasumasa Onoe, Michael Zhang, Eunsol Choi, and Greg Durrett. Entity cloze by date: What LMs
949 know about unseen entities. In Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir
950 Meza Ruiz (eds.), *Findings of the Association for Computational Linguistics: NAACL 2022*,
951 pp. 693–702, Seattle, United States, July 2022. Association for Computational Linguistics.
952 doi: 10.18653/v1/2022.findings-naacl.52. URL [https://aclanthology.org/2022.findings-naacl.52/](https://aclanthology.org/2022.findings-naacl.52).

953

954 Yasumasa Onoe, Michael Zhang, Shankar Padmanabhan, Greg Durrett, and Eunsol Choi. Can
955 LMs learn new entities from descriptions? challenges in propagating injected knowledge. In
956 Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual
957 Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5469–
958 5485, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/
959 2023.acl-long.300. URL <https://aclanthology.org/2023.acl-long.300/>.

960

961 Riccardo Orlando, Pere-Lluís Huguet Cabot, Edoardo Barba, and Roberto Navigli. ReLiK: Re-
962 trieve and LinK, fast and accurate entity linking and relation extraction on an academic bud-
963 get. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Associa-
964 tion for Computational Linguistics: ACL 2024*, pp. 14114–14132, Bangkok, Thailand, August
965 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.839. URL
<https://aclanthology.org/2024.findings-acl.839/>.

966

967 Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. Unifying large
968 language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data
969 Engineering*, 2024.

970

971 Heiko Paulheim. Knowledge graph refinement: A survey of approaches and evaluation methods.
Semantic web, 8(3):489–508, 2016. doi: 10.3233/SW-160218. URL <https://journals.sagepub.com/doi/abs/10.3233/SW-160218>.

972 Flavio Petruzzellis, Cristina Cornelio, and Pietro Lio. Hierarchical planning for complex tasks with
973 knowledge graph-rag and symbolic verification. In *Forty-second International Conference on*
974 *Machine Learning*, 2025.

975 Axel Polleres, Romana Pernisch, Angela Bonifati, Daniele Dell’Aglio, Daniil Dobriy, Stefania
976 Dumbrava, Lorena Etcheverry, Nicolas Ferranti, Katja Hose, Ernesto Jiménez-Ruiz, et al. How
977 does knowledge evolve in open knowledge graphs? *Transactions on Graph Data and Knowledge*,
978 1(1):11–1, 2023.

979 Simon Razniewski, Hiba Arnaout, Shrestha Ghosh, and Fabian Suchanek. Completeness, recall,
980 and negation in open-world knowledge bases: A survey. *ACM Computing Surveys*, 56(6):1–42,
981 2024. doi: 10.1145/3639563. URL <https://doi.org/10.1145/3639563>.

982 Ridho Reinanda, Edgar Meij, Maarten de Rijke, et al. Knowledge graphs: An information retrieval
983 perspective. *Foundations and Trends® in Information Retrieval*, 14(4):289–444, 2020.

984 Sebastian Riedel, Limin Yao, and Andrew McCallum. Modeling relations and their mentions
985 without labeled text. In *Proceedings of the 2010 European Conference on Machine Learn-*
986 *ing and Knowledge Discovery in Databases*, pp. 148–163, 2010. doi: https://doi.org/10.1007/978-3-642-15939-8_10. URL https://link.springer.com/chapter/10.1007/978-3-642-15939-8_10.

987 Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Nandana Mihindukulasooriya, Owen Corne, and
988 Alfio Massimiliano Gliozzo. Knowgl: Knowledge generation and linking from text. In *Pro-*
989 *ceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 16476–16478, 2023.
990 doi: <https://doi.org/10.1609/aaai.v37i13.27084>. URL <https://ojs.aaai.org/index.php/AAAI/article/view/27084>.

991 Dan Roth and Wen-tau Yih. A linear programming formulation for global inference in natural
992 language tasks. Technical report, Illinois Univ at Urbana-Champaign Dept of Computer Science,
993 2004. URL <https://aclanthology.org/W04-2401.pdf>.

994 Tara Safavi and Danai Koutra. CoDEX: A Comprehensive Knowledge Graph Completion Bench-
995 mark. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020*
996 *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 8328–8350,
997 Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.
998 *emnlp-main.669*. URL <https://aclanthology.org/2020.emnlp-main.669/>.

999 Alessandro Seganti, Klaudia Firlkag, Helena Skowronska, Michal Satlawa, and Piotr An-
1000 druszkiewicz. Multilingual entity and relation extraction dataset and model. In *Proceed-
1001 ings of the 16th conference of the european chapter of the association for computational lin-
1002 guistics: Main volume*, pp. 1946–1955, 2021. doi: 10.18653/v1/2021.eacl-main.166. URL
1003 <https://aclanthology.org/2021.eacl-main.166/>.

1004 Soumya Sharma, Tapas Nayak, Arusarka Bose, Ajay Kumar Meena, Koustuv Dasgupta, Niloy Gan-
1005 guly, and Pawan Goyal. FinRED: A dataset for relation extraction in financial domain. In *Com-
1006 panion Proceedings of the Web Conference 2022*, pp. 595–597, 2022. doi: 10.48550/ARXIV.
1007 2306.03736. URL <https://doi.org/10.48550/arXiv.2306.03736>.

1008 Tong Shen, Fu Zhang, and Jingwei Cheng. A comprehensive overview of knowledge graph com-
1009 pletion. *Knowledge-Based Systems*, 255:109597, 2022. doi: <https://doi.org/10.1016/j.knosys.2022.109597>. URL <https://www.sciencedirect.com/science/article/pii/S095070512200805X>.

1010 Zhiyi Song, Ann Bies, Stephanie Strassel, Tom Riese, Justin Mott, Joe Ellis, Jonathan Wright, Seth
1011 Kulick, Neville Ryant, and Xiaoyi Ma. From light to rich ere: annotation of entities, relations, and
1012 events. In *Proceedings of the the 3rd Workshop on EVENTS: Definition, Detection, Coreference,
1013 and Representation*, pp. 89–98, 2015. URL <https://aclanthology.org/W15-0812.pdf>.

1014 George Stoica, Emmanouil Antonios Platanios, and Barnabás Póczos. Re-TACRED: Address-
1015 ing shortcomings of the tacred dataset. In *Proceedings of the AAAI conference on artificial
1016 intelligence*, volume 35, pp. 13843–13850, 2021. doi: 10.1609/AAAI.V35I15.17631. URL
1017 <https://doi.org/10.1609/aaai.v35i15.17631>.

1026 Budhitama Subagdja, D Shanthoshigaa, Zhaoxia Wang, and Ah-Hwee Tan. Machine learning for
1027 refining knowledge graphs: A survey. *ACM Computing Surveys*, 56(6):1–38, 2024. doi: 10.1145/
1028 3640313. URL <https://doi.org/10.1145/3640313>.

1029 Jiahuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M
1030 Ni, Heung-Yeung Shum, and Jian Guo. Think-on-graph: Deep and responsible reasoning of
1031 large language model on knowledge graph. In *The Twelfth International Conference on Learn-
1032 ing Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*, 2024. URL <https://openreview.net/forum?id=nnVO1PvbTv>.

1033 TAC-KBP. TAC-KBP home page, 2022. URL <https://tac.nist.gov/tracks/index.html>.

1034 Qingyu Tan, Lu Xu, Lidong Bing, Hwee Tou Ng, and Sharifah Mahani Aljunied. Revisit-
1035 ing docRED-addressing the false negative problem in relation extraction. In *Proceedings of
1036 the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 8472–8487,
1037 2022. doi: 10.18653/v1/2022.emnlp-main.580. URL [https://aclanthology.org/2022.emnlp-main.580/](https://aclanthology.org/2022.emnlp-main.580).

1038 Kristina Toutanova and Danqi Chen. Observed versus latent features for knowledge base and text
1039 inference. In *Proceedings of the 3rd workshop on continuous vector space models and their
1040 compositionality*, pp. 57–66, 2015. URL <https://aclanthology.org/W15-4007.pdf>.

1041 Denny Vrandečić and Markus Krötzsch. Wikidata: a free collaborative knowledgebase. *Commu-
1042 nications of the ACM*, 57(10):78–85, 2014. doi: 10.1145/2629489. URL <https://dl.acm.org/doi/fullHtml/10.1145/2629489>.

1043 Tu Vu, Mohit Iyyer, Xuezhi Wang, Noah Constant, Jerry Wei, Jason Wei, Chris Tar, Yun-Hsuan
1044 Sung, Denny Zhou, Quoc Le, and Thang Luong. FreshLLMs: Refreshing large language mod-
1045 els with search engine augmentation. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar
1046 (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 13697–13720,
1047 Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/
1048 v1/2024.findings-acl.813. URL [https://aclanthology.org/2024.findings-acl.813/](https://aclanthology.org/2024.findings-acl.813).

1049 Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. ACE 2005 multilingual
1050 training corpus. *Linguistic Data Consortium, Philadelphia*, 57, 2006. doi: <https://doi.org/10.35111/mwxc-vh88>. URL <https://doi.org/10.35111/mwxc-vh88>.

1051 Liang Wang, Wei Zhao, Zhuoyu Wei, and Jingming Liu. SimKGC: Simple contrastive knowl-
1052 edge graph completion with pre-trained language models. In Smaranda Muresan, Preslav Nakov,
1053 and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for
1054 Computational Linguistics (Volume 1: Long Papers)*, pp. 4281–4294, Dublin, Ireland, May
1055 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.295. URL
1056 [https://aclanthology.org/2022.acl-long.295/](https://aclanthology.org/2022.acl-long.295).

1057 Shijie Wang, Wenqi Fan, Yue Feng, Lin Shanru, Xinyu Ma, Shuaiqiang Wang, and Dawei Yin.
1058 Knowledge graph retrieval-augmented generation for llm-based recommendation. In *Proceedings
1059 of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long
1060 Papers)*, pp. 27152–27168, 2025.

1061 Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhiyuan Liu, Juanzi Li, and Jian Tang. KEPLER: A
1062 unified model for knowledge embedding and pre-trained language representation. *Transactions of
1063 the Association for Computational Linguistics*, 9:176–194, 03 2021. doi: 10.1162/tacl_a_00360.
1064 URL https://doi.org/10.1162/tacl_a_00360.

1065 Yu Wang, Nedim Lipka, Ryan A Rossi, Alexa Siu, Ruiyi Zhang, and Tyler Derr. Knowledge graph
1066 prompting for multi-document question answering. In *Proceedings of the AAAI Conference on
1067 Artificial Intelligence*, volume 38, pp. 19206–19214, 2024.

1068 Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-
1069 Ziv, Neel Jain, Khalid Saifullah, Siddartha Naidu, et al. Livebench: A challenging, contamina-
1070 tion-free llm benchmark. *arXiv preprint arXiv:2406.19314*, 2024. URL <https://openreview.net/forum?id=sKYHBTaxVa>.

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080 Tongtong Wu, Linhao Luo, Yuan-Fang Li, Shirui Pan, Thuy-Trang Vu, and Gholamreza Haffari.
1081 Continual learning for large language models: A survey. *arXiv preprint arXiv:2402.01364*, 2024a.
1082 URL <https://arxiv.org/pdf/2402.01364.pdf>.

1083 Xiaobao Wu, Liangming Pan, William Yang Wang, and Anh Tuan Luu. AKEW: Assessing knowl-
1084 edge editing in the wild. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Pro-
1085 ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp.
1086 15118–15133, Miami, Florida, USA, November 2024b. Association for Computational Linguis-
1087 tics. doi: 10.18653/v1/2024.emnlp-main.843. URL <https://aclanthology.org/2024.emnlp-main.843>.

1088
1089 Xiaobao Wu, Liangming Pan, Yuxi Xie, Ruiwen Zhou, Shuai Zhao, Yubo Ma, Mingzhe Du, Rui
1090 Mao, Anh Tuan Luu, and William Yang Wang. Antileak-bench: Preventing data contamination
1091 by automatically constructing benchmarks with updated real-world knowledge. *arXiv preprint
1092 arXiv:2412.13670*, 2024c. URL <https://arxiv.org/pdf/2412.13670.pdf>.

1093
1094 Rui Xing, Jie Luo, and Tengwei Song. BioRel: towards large-scale biomedical relation extraction.
1095 *BMC bioinformatics*, 21:1–13, 2020. doi: 10.1186/S12859-020-03889-5. URL <https://doi.org/10.1186/s12859-020-03889-5>.

1096
1097 Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. One-shot relational
1098 learning for knowledge graphs. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi
1099 Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language
1100 Processing*, pp. 1980–1990, Brussels, Belgium, October–November 2018. Association for Com-
1101 putational Linguistics. doi: 10.18653/v1/D18-1223. URL <https://aclanthology.org/D18-1223>.

1102
1103 Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng
1104 Zheng, Yang Wang, and Enhong Chen. Large language models for generative information
1105 extraction: A survey. *Frontiers of Computer Science*, 18(6):186357, 2024. doi: 10.1007/
1106 S11704-024-40555-Y. URL <https://doi.org/10.1007/s11704-024-40555-y>.

1107
1108 Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie
1109 Zhou, and Maosong Sun. DocRED: A large-scale document-level relation extraction dataset. In
1110 *Proceedings of the 2019 Annual Meeting of the Association for Computational Linguistics (ACL
1111 2019)*, pp. 764–777, 2019. URL <https://aclanthology.org/P19-1074>.

1112
1113 Dian Yu, Kai Sun, Claire Cardie, and Dong Yu. Dialogue-based relation extraction. In *Proceedings
1114 of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4927–4940,
1115 2020. doi: 10.18653/v1/2020.acl-main.444. URL <https://aclanthology.org/2020.acl-main.444>.

1116
1117 Klim Zaporojets, Johannes Deleu, Chris Develder, and Thomas Demeester. DWIE: An entity-centric
1118 dataset for multi-task document-level information extraction. *Information Processing & Manage-
1119 ment*, 58(4):102563, 2021. URL <https://doi.org/10.1016/j.ipm.2021.102563>.

1120
1121 Urchade Zaratiana, Nadi Tomeh, Pierre Holat, and Thierry Charnois. An autoregressive text-to-
1122 graph framework for joint entity and relation extraction. In *Proceedings of the AAAI Conference on
1123 Artificial Intelligence*, volume 38, pp. 19477–19487, 2024. doi: 10.1609/AAAI.V38I17.29919.
1124 URL <https://doi.org/10.1609/aaai.v38i17.29919>.

1125
1126 Bowen Zhang and Harold Soh. Extract, define, canonicalize: An LLM-based framework for knowl-
1127 edge graph construction. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Pro-
1128 ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp.
1129 9820–9836, Miami, Florida, USA, November 2024. Association for Computational Linguis-
1130 tics. doi: 10.18653/v1/2024.emnlp-main.548. URL <https://aclanthology.org/2024.emnlp-main.548>.

1131
1132 Dongxu Zhang, Sunil Mohan, Michaela Torkar, and Andrew Mccallum. A distant supervision cor-
1133 pus for extracting biomedical relationships between chemicals, diseases and genes. In *Proceed-
1134 ings of the Thirteenth Language Resources and Evaluation Conference*, pp. 1073–1082, 2022.
1135 URL <https://aclanthology.org/2022.lrec-1.116>.

1134 Jin-Cheng Zhang, Azlan Mohd Zain, Kai-Qing Zhou, Xi Chen, and Ren-Min Zhang. A review of
1135 recommender systems based on knowledge graph embedding. *Expert Systems With Applications*,
1136 250:123876, 2024.

1137 Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. Position-
1138 aware attention and supervised data improve slot filling. In Martha Palmer, Rebecca Hwa, and
1139 Sebastian Riedel (eds.), *Proceedings of the 2017 Conference on Empirical Methods in Natu-
1140 ral Language Processing*, pp. 35–45, Copenhagen, Denmark, September 2017. Association for
1141 Computational Linguistics. doi: 10.18653/v1/D17-1004. URL <https://aclanthology.org/D17-1004/>.

1142

1143 Zikang Zhang, Wangjie You, Tianci Wu, Xinrui Wang, Juntao Li, and Min Zhang. A survey of
1144 generative information extraction. In *Proceedings of the 31st International Conference on Com-
1145 putational Linguistics*, pp. 4840–4870, 2025. URL <https://aclanthology.org/2025.coling-main.324/>.

1146

1147 Bowen Zhao, Zander Brumbaugh, Yizhong Wang, Hannaneh Hajishirzi, and Noah Smith. Set the
1148 clock: Temporal alignment of pretrained language models. In Lun-Wei Ku, Andre Martins, and
1149 Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*,
1150 pp. 15015–15040, Bangkok, Thailand, August 2024a. Association for Computational Linguis-
1151 tics. doi: 10.18653/v1/2024.findings-acl.892. URL <https://aclanthology.org/2024.findings-acl.892/>.

1152

1153 Xiaoyan Zhao, Yang Deng, Min Yang, Lingzhi Wang, Rui Zhang, Hong Cheng, Wai Lam, Ying
1154 Shen, and Rui Feng Xu. A comprehensive survey on relation extraction: Recent advances and new
1155 frontiers. *ACM Computing Surveys*, 56(11):1–39, 2024b.

1156

1157 Xiaohan Zou. A survey on application of knowledge graph. In *Journal of Physics: Conference
1158 Series*, volume 1487, pp. 012016. IOP Publishing, 2020.

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188 A EXTENDED RELATED WORK
1189

1190 This appendix provides an expanded discussion of related work, offering additional context, com-
1191 parisons, and references beyond those included in the main text.
1192

1193 A.1 COMPARISON OF EMERGE WITH EXISTING INFORMATION EXTRACTION BENCHMARKS
1194

1195 Table 4 presents a detailed comparison of EMERGE with existing information extraction (IE) bench-
1196 mark datasets across the following key criteria:
1197

- **Evolution:** indicates whether the dataset captures the natural evolution of knowledge in knowledge graph (*KG*) and textual (*Text*) sources.
- **Text-to-KG integration:** extent to which information extraction annotations are integrated with knowledge in a KG, broken down in:
 - **KG Link:** indicates whether the annotated entities in the triples are linked to a KG, supporting thus *entity linking* task.
 - **X-Triples:** presence of triples aligned with facts already in a KG (*X-Triples* TKGU operation; Section 3).
 - **E-Triples:** whether a dataset can be used to extract triples from text that connect existing entities in a KG (*E-Triples* TKGU type; Section 3).
 - **EE-Triples:** coverage of triples involving emerging (non-existing) entities in a KG (*EE-Triples* TKGU; Section 3).
 - **EE-KG-Triples:** availability of annotations linking emerging entities in text to other entities in KG not mentioned in text (*EE-KG-Triples* TKGU; Section 3).
 - **D-Triples:** inclusion of annotations that mark deprecation of existing KG triples based on information in textual passage (*D-Triples* TKGU; Section 3).

1215 From Table 4, we observe that, to the best of our knowledge, none of the existing IE datasets support
1216 information extraction in a realistic knowledge evolution setting, where knowledge evolves simul-
1217 taneously in both KG and textual sources (columns *Evolution-KG* and *Evolution-Text* in the table).
1218 Moreover, a number of datasets, such as TACRED (Zhang et al., 2017), BC5CDR (Li et al., 2016),
1219 DDI (Herrero-Zazo et al., 2013), and DWIE (Zaporojets et al., 2021), include *entity linking* to a
1220 KG, but are not accompanied by an actual KG, and their extracted relations do not align directly
1221 with the relations defined in a KG schema. Finally, although E-Triples and EE-Triples operations
1222 are nominally supported in some of the compared datasets, they do not capture genuinely emerging
1223 knowledge; instead, they rely on random subsampling of triples to approximate TKGU operations.
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

1242
1243 Table 4: Overview of major information extraction datasets from the past three decades across
1244 various domains, compared to our EMERGE dataset.

1245 1246 1247 1248 1249 1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292 1293 1294 1295 1296 1297 1298 1299 1300 1301 1302 1303 1304 1305 1306 1307 1308 1309 1310 1311 1312 1313 1314 1315 1316 1317 1318 1319 1320 1321 1322 1323 1324 1325 1326 1327 1328 1329 1330 1331 1332 1333 1334 1335 1336 1337 1338 1339 1340 1341 1342 1343 1344 1345 1346 1347 1348 1349 1350 1351 1352 1353 1354 1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1400 1401 1402 1403 1404 1405 1406 1407 1408 1409 1410 1411 1412 1413 1414 1415 1416 1417 1418 1419 1420 1421 1422 1423 1424 1425 1426 1427 1428 1429 1430 1431 1432 1433 1434 1435 1436 1437 1438 1439 1440 1441 1442 1443 1444 1445 1446 1447 1448 1449 1450 1451 1452 1453 1454 1455 1456 1457 1458 1459 1460 1461 1462 1463 1464 1465 1466 1467 1468 1469 1470 1471 1472 1473 1474 1475 1476 1477 1478 1479 1480 1481 1482 1483 1484 1485 1486 1487 1488 1489 1490 1491 1492 1493 1494 1495 1496 1497 1498 1499 1500 1501 1502 1503 1504 1505 1506 1507 1508 1509 1510 1511 1512 1513 1514 1515 1516 1517 1518 1519 1520 1521 1522 1523 1524 1525 1526 1527 1528 1529 1530 1531 1532 1533 1534 1535 1536 1537 1538 1539 1540 1541 1542 1543 1544 1545 1546 1547 1548 1549 1550 1551 1552 1553 1554 1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565 1566 1567 1568 1569 1570 1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619 1620 1621 1622 1623 1624 1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 1674 1675 1676 1677 1678 1679 1680 1681 1682 1683 1684 1685 1686 1687 1688 1689 1690 1691 1692 1693 1694 1695 1696 1697 1698 1699 1700 1701 1702 1703 1704 1705 1706 1707 1708 1709 1710 1711 1712 1713 1714 1715 1716 1717 1718 1719 1720 1721 1722 1723 1724 1725 1726 1727 1728 1729 1730 1731 1732 1733 1734 1735 1736 1737 1738 1739 1740 1741 1742 1743 1744 1745 1746 1747 1748 1749 1750 1751 1752 1753 1754 1755 1756 1757 1758 1759 1760 1761 1762 1763 1764 1765 1766 1767 1768 1769 1770 1771 1772 1773 1774 1775 1776 1777 1778 1779 1780 1781 1782 1783 1784 1785 1786 1787 1788 1789 1790 1791 1792 1793 1794 1795 1796 1797 1798 1799 1800 1801 1802 1803 1804 1805 1806 1807 1808 1809 1810 1811 1812 1813 1814 1815 1816 1817 1818 1819 1820 1821 1822 1823 1824 1825 1826 1827 1828 1829 1830 1831 1832 1833 1834 1835 1836 1837 1838 1839 1840 1841 1842 1843 1844 1845 1846 1847 1848 1849 1850 1851 1852 1853 1854 1855 1856 1857 1858 1859 1860 1861 1862 1863 1864 1865 1866 1867 1868 1869 1870 1871 1872 1873 1874 1875 1876 1877 1878 1879 1880 1881 1882 1883 1884 1885 1886 1887 1888 1889 1890 1891 1892 1893 1894 1895 1896 1897 1898 1899 1900 1901 1902 1903 1904 1905 1906 1907 1908 1909 1910 1911 1912 1913 1914 1915 1916 1917 1918 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1930 1931 1932 1933 1934 1935 1936 1937 1938 1939 1940 1941 1942 1943 1944 1945 1946 1947 1948 1949 1950 1951 1952 1953 1954 1955 1956 1957 1958 1959 1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980 1981 1982 1983 1984 1985 1986 1987 1988 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023 2024 2025 2026 2027 2028 2029 2030 2031 2032 2033 2034 2035 2036 2037 2038 2039 2040 2041 2042 2043 2044 2045 2046 2047 2048 2049 2050 2051 2052 2053 2054 2055 2056 2057 2058 2059 2060 2061 2062 2063 2064 2065 2066 2067 2068 2069 2070 2071 2072 2073 2074 2075 2076 2077 2078 2079 2080 2081 2082 2083 2084 2085 2086 2087 2088 2089 2090 2091 2092 2093 2094 2095 2096 2097 2098 2099 20100 20101 20102 20103 20104 20105 20106 20107 20108 20109 20110 20111 20112 20113 20114 20115 20116 20117 20118 20119 20120 20121 20122 20123 20124 20125 20126 20127 20128 20129 20130 20131 20132 20133 20134 20135 20136 20137 20138 20139 20140 20141 20142 20143 20144 20145 20146 20147 20148 20149 20150 20151 20152 20153 20154 20155 20156 20157 20158 20159 20160 20161 20162 20163 20164 20165 20166 20167 20168 20169 20170 20171 20172 20173 20174 20175 20176 20177 20178 20179 20180 20181 20182 20183 20184 20185 20186 20187 20188 20189 20190 20191 20192 20193 20194 20195 20196 20197 20198 20199 20200 20201 20202 20203 20204 20205 20206 20207 20208 20209 20210 20211 20212 20213 20214 20215 20216 20217 20218 20219 20220 20221 20222 20223 20224 20225 20226 20227 20228 20229 202210 202211 202212 202213 202214 202215 202216 202217 202218 202219 202220 202221 202222 202223 202224 202225 202226 202227 202228 202229 202230 202231 202232 202233 202234 202235 202236 202237 202238 202239 202240 202241 202242 202243 202244 202245 202246 202247 202248 202249 202250 202251 202252 202253 202254 202255 202256 202257 202258 202259 202260 202261 202262 202263 202264 202265 202266 202267 202268 202269 202270 202271 202272 202273 202274 202275 202276 202277 202278 202279 202280 202281 202282 202283 202284 202285 202286 202287 202288 202289 202290 202291 202292 202293 202294 202295 202296 202297 202298 202299 2022100 2022101 2022102 2022103 2022104 2022105 2022106 2022107 2022108 2022109 2022110 2022111 2022112 2022113 2022114 2022115 2022116 2022117 2022118 2022119 2022120 2022121 2022122 2022123 2022124 2022125 2022126 2022127 2022128 2022129 2022130 2022131 2022132 2022133 2022134 2022135 2022136 2022137 2022138 2022139 2022140 2022141 2022142 2022143 2022144 2022145 2022146 2022147 2022148 2022149 2022150 2022151 2022152 2022153 2022154 2022155 2022156 2022157 2022158 2022159 2022160 2022161 2022162 2022163 2022164 2022165 2022166 2022167 2022168 2022169 2022170 2022171 2022172 2022173 2022174 2022175 2022176 2022177 2022178 2022179 2022180 2022181 2022182 2022183 2022184 2022185 2022186 2022187 2022188 2022189 2022190 2022191 2022192 2022193 2022194 2022195 2022196 2022197 2022198 2022199 2022200 2022201 2022202 2022203 2022204 2022205 2022206 2022207 2022208 2022209 2022210 2022211 2022212 2022213 2022214 2022215 2022216 2022217 2022218 2022219 2022220 2022221 2022222 2022223 2022224 2022225 2022226 2022227 2022228 2022229 20222210 20222211 20222212 20222213 20222214 20222215 20222216 20222217 20222218 20222219 20222220 20222221 20222222 20222223 20222224 20222225 20222226 20222227 20222228 20222229 202222210 202222211 202222212 202222213 202222214 202222215 202222216 202222217 202222218 202222219 202222220 202222221 202222222 202222223 202222224 202222225 202222226 202222227 202222228 202222229 2022222210 2022222211 2022222212 2022222213 2022222214 2022222215 2022222216 2022222217 2022222218 2022222219 2022222220 2022222221 2022222222 2022222223 2022222224 2022222225 2022222226 2022222227 2022222228 2022222229 20222222210 20222222211 20222222212 20222222213 20222222214 20222222215 20222222216 20222222217 20222222218 20222222219 20222222220 20222222221 20222222222 20222222223 20222222224 20222222225 20222222226 20222222227 20222222228 20222222229 202222222210 202222222211 202222222212 202222222213 202222222214 202222222215 202222222216 202222222217 202222222218 202222222219 202222222220 202222222221 202222222222 202222222223 202222222224 202222222225 202222222226 202222222227 202222222228 202222222229 2022222222210 2022222222211 2022222222212 2022222222213 2022222222214 2022222222215 2022222222216 2022222222217 2022222222218 2022222222219 2022222222220 2022222222221 2022222222222 2022222222223 2022222222224 2022222222225 2022222222226 2022222222227 2022222222228 2022222222229 20222222222210 20222222222211 20222222222212 20222222222213 20222222222214 20222222222215 20222222222216 20222222222217 20222222222218 20222222222219 20222222222220 20222222222221 20222222222222 20222222222223 20222222222224 20222222222225 20222222222226 20222222222227 20222222222228 20222222222229 202222222222210 202222222222211 202222222222212 202222222222213 202222222222214 202222222222215 202222222222216 202222222222217 202222222222218 202222222222219 202222222222220 202222222222221 202222222222222 202222222222223 202222222222224 202222222222225 202222222222226 202222222222227 202222222222228 202222222222229 2022222222222210 2022222222222211 2022222222222212 2022222222222213 2022222222222214 2022222222222215 2022222222222216 2022222222222217 2022222222222218 2022222222222219 2022222222222220 2022222222222221 2022222222222222 2022222222222223 2022222222222224 2022222222222225 2022222222222226 2022222222222227 2022222222222228 2022222222222229 20222222222222210 20222222222222211 20222222222222212 20222222222222213 20222222222222214 20222222222222215 20222222222222216 20222222222222217 20222222222222218 20222222222222219 20222222222222220 20222222222222221 20222222222222222 20222222222222223 20222222222222224 20222222222222225 20222222222222226 20222222222222227 20222222222222228 20222222222222229 202222222222222210 202222222222222211 202222222222222212 202222222222222213 202222222222222214 202222222222222215 202222222222222216 202222222222222217 202222222222222218 202222222222222219 202222222222222220 202222222222222221 202222222222222222 202222222222222223 202222222222222224 202222222222222225 202222222222222226 202222222222222227 202222222222222228 202222222222222229 2022222222222222210 2022222222222222211 2022

1296 where the main task is to predict missing relations between entities. Thus, in their work, Bordes
1297 et al. (2013) introduced the WN18 and FB15k datasets. These datasets are derived from WordNet
1298 (Miller, 1995) and Freebase (Bollacker et al., 2008) respectively and capture the relations between
1299 entities. Later work (Toutanova & Chen, 2015; Dettmers et al., 2018) modified WN18 and FB15k
1300 datasets to eliminate redundant relations and train-test leakage, leading to the release of WN18RR
1301 and FB15K-237 datasets. More recently, a much larger Wikidata5M Wang et al. (2021) was re-
1302 leased and contains ~ 5 million entities and ~ 20 million triples. Other widely used text-based KGC
1303 datasets are Wiki/NELL-One (Xiong et al., 2018), FB15K-237N (Lv et al., 2022), CoDEX (Safavi &
1304 Koutra, 2020), YAGO3-10 (Mahdisoltani et al., 2014) and LiterallyWikidata (Gesese et al., 2021).
1305 While these datasets enable models to incorporate textual information as node features (Daza et al.,
1306 2021; Wang et al., 2022), they remain static and do not capture the evolving nature of knowledge
1307 within KGs. Moreover, the KG triples in these datasets are not linked to textual sources that repre-
1308 sent their information. To address this gap, our dataset captures the evolution of knowledge in the
1309 Wikidata KG and links KG updates to textual evidence from passages in Wikipedia pages. In doing
1310 so, it also creates opportunities to integrate ideas from KG completion, such as enforcing structural
1311 consistency, into text-driven information extraction (see next paragraph), thereby bridging two lines
1312 of work that are typically treated separately.

1313 **Information extraction (IE).** To evaluate the ability of models to extract structured knowledge,
1314 researchers have developed IE datasets by annotating entity relations. MUC-7 (Chinchor & Marsh,
1315 1998) introduced three relation types, with later datasets expanding in size, relation diversity, or both.
1316 Notable examples include CoNLL04 (Roth & Yih, 2004), ACE 2005 (Walker et al., 2006), ERE
1317 (Aguilar et al., 2014; Song et al., 2015), BC5CDR (Li et al., 2016), TACRED (Zhang et al., 2017),
1318 SciERC (Luan et al., 2018), SemEval-2010 (Hendrickx et al., 2010), SemEval-2017 (Augenstein
1319 et al., 2017), DWIE (Zaporojets et al., 2021) and BioRED (Luo et al., 2022), among others. Other
1320 datasets, such as NYT (Riedel et al., 2010), explicitly linked KG triples to textual snippets using
1321 distant supervision. Similarly, but on a larger scale, Gabrilovich et al. (2013) introduced FACC1 by
1322 aligning ClueWeb12 documents with Freebase entity mention annotations. In parallel, the TAC-KBP
1323 challenges (Ji et al., 2010; TAC-KBP, 2022) (2009 – 2020) produced proprietary manually annotated
1324 datasets for knowledge base population tasks such as slot filling and entity linking. More recently,
1325 these resources have been extended with a variety of datasets that map textual knowledge to KG
1326 literals, such as LiterallyWikidata (Gesese et al., 2021), and KG triples, such as WebNLG (Gardent
1327 et al., 2017), KELM (Agarwal et al., 2021), FewRel (Han et al., 2018), DocRED (Yao et al., 2019),
1328 Wiki/GEO-NRE (Distiawan et al., 2019), BioRel (Xing et al., 2020), T-REX (Elsahar et al., 2019)
1329 and REBEL (Cabot & Navigli, 2021). While these datasets connect textual knowledge to KG triples,
1330 they do not account for the operations required to update a KG as new information emerges in text.
1331 Our work addresses this gap by linking new textual knowledge to the specific update operations (see
1332 Section 3) on a KG snapshot. Methodologically, existing state-of-the-art IE methods (see Table 1)
1333 provide a natural starting point for tackling TKGU operations, as they extract structured knowledge
1334 from text and already cover some of these operations. However, they remain largely oblivious to the
1335 existing KG structure and require extensions to integrate emerging textual content into the KG, such
1336 as deprecating outdated triples, adding new entities, and enforcing structural consistency based on
1337 how entities and relations are used within the KG.

1338 **Continual learning with emerging knowledge.** Over the last few years, there has been a grow-
1339 ing interest in developing datasets aimed at probing models on emerging knowledge. Datasets like
1340 ECBD (Onoe et al., 2022), TemporalWiki (Jang et al., 2022a), TempLAMA (Dhingra et al., 2022),
1341 DynamicTempLAMA (Margatina et al., 2023), Updated and New LAMA (Jang et al., 2022b) were
1342 proposed to evaluate LLMs on slot-filling tasks using up-to-date knowledge. More recently, this
1343 line of work has expanded to question answering on emerging knowledge, with datasets such as
1344 StreamingQA (Liska et al., 2022), FreshQA (Vu et al., 2024), EvolvingQA (Kim et al., 2024b),
1345 RealtimeQA (Kasai et al., 2024), DynamicER (Kim et al., 2024a), GrowOVER (Ko et al., 2024),
1346 ERASE (Li et al., 2025a), Wiki-Update (Wu et al., 2024b), AntiLeak-Bench (Wu et al., 2024c), and
1347 LiveBench (White et al., 2024). However, existing datasets do not evaluate models on dynamically
1348 updating large-scale KGs while grounding changes in textual evidence. This setting requires mod-
1349 els to be aware of changes in continually evolving KG schema and emerging knowledge in textual
sources. To address this, we introduce EMERGE, a dataset that links emerging textual knowledge
to updates in a time-evolving Wikidata KG with 37 million edges.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

KG versioning. Our work is also related to KG versioning (Jin et al., 2025; Alam et al., 2024; Hofer et al., 2024). Similar to prior work in this area, we construct a KG in which each edge is annotated with its temporal span, capturing both its addition and deprecation history. This enables efficient extraction of KG snapshots and deltas (see Section 4.1). However, in contrast to KG versioning approaches that focus solely on maintaining temporal KG states, we use these versions as an intermediate step to build a dataset where KG updates are linked to the textual evidence in the corresponding Wikipedia passages. As a result, unlike existing IE datasets in which annotated triples are selected independently of KG evolution, the triples in EMERGE reflect the natural progression of facts in a real-world KG, making the dataset highly practical and grounded in authentic knowledge change.

B METRICS

B.1 RECALL

We use recall, which measures the fraction of correctly predicted ground truth triples and is defined as follows:

$$\text{Recall} = \frac{|\mathcal{T}_D \cap \mathcal{T}'_D|}{|\mathcal{T}'_D|},$$

where \mathcal{T}_D is a set of predicted triples and \mathcal{T}'_D is the set of ground truth triples.

B.2 COMPLETENESS

The completeness metric (Jiang et al., 2024) can be formalized as follows:

$$c(\mathcal{T}'_D, \mathcal{T}_D) = \frac{|\{\tau \in \mathcal{T}'_D | \exists \tau \in \mathcal{T}_D, \text{sim}(\tau, \tau') \geq \phi\}|}{|\mathcal{T}'_D|},$$

where \mathcal{T}'_D is the set of ground truth, and \mathcal{T}_D the set of predicted triples. $\text{sim}(\tau, \tau') = \text{CosSim}(\text{emb}(\tau), \text{emb}(\tau'))$. We use SentenceTransformer('all-mpnet-base-v2') to calculate the embeddings emb . We set the threshold ϕ to 0.9, which, based on our observations, provides accurate similarity matching.

C QUALITY CONTROL

In this section, we describe how LLMs are used to automatically filter out triples that cannot be derived from textual passages (Section C.1). We also detail the human annotation process used to validate the resulting LLM-generated annotations (Section C.2).

C.1 QUALITY CONTROL PROMPTS AND EXAMPLES

We use two different prompts to filter out triples that cannot be inferred from a textual passage. The first is an *assertion prompt* (see Section C.1.1) applied to validate *X-Triples*, *E-Triples*, *EE-Triples*, and *EE-KG-Triples* as defined in Section 3. The goal of this prompt is to verify whether a triple can be directly or indirectly derived from the text. The second prompt is a *deprecation prompt* (see Section C.1.2), and is used to validate the deprecation of triples involved in *D-Triples* TKGU operation.

C.1.1 TRIPLE ASSERTION PROMPT

The following is the structure of the prompt used to assert that the *X-Triples*, *E-Triples*, *EE-Triples*, and *EE-KG-Triples* TKGU operations can be derived from the information in textual passages. The placeholder <TEXT> is replaced by the textual passage, and <TRIPLES_LIST> by a list of triples.

You are given the following text:

<TEXT>

1404
1405 Can the following triples be directly or indirectly (the text
1406 provides some hints) inferred from the text? Use common sense
1407 but not knowledge that cannot be inferred from the text above.

1408 <TRIPLES_LIST>
1409

1410 Write a numbered list with the triples above, where each of the
1411 triples is followed by YES if the triple is represented in the
1412 text, and NO otherwise. Follow your 'YES' or 'NO' answer for
1413 each triple with a brief explanation.

1414 The following is a concrete example of the *triple assertion* prompt outlined above:
1415

1416 You are given the following text:
1417

1418 After the war, Abraham H. Albertson returned to complete his
1419 education and graduated, with scholarship assistance, from the
1420 Columbia School of Architecture with a Ph.B in 1895.
1421

1422 Can the following triples be directly or indirectly (the text
1423 provides some hints) inferred from the text? Use common sense
1424 but not knowledge that cannot be inferred from the text above.
1425

1426 1. [Abraham H. Albertson (Q80520646); educated at (P69);
1427 Columbia Graduate School of Architecture, Planning and
1428 Preservation (Q5149701)]
1429 2. [Abraham H. Albertson (Q80520646); occupation (P106);
1430 architect (Q42973)]
1431 3. [Abraham H. Albertson (Q80520646); country of
1432 citizenship (P27); United States of America (Q30)]
1433 4. [Abraham H. Albertson (Q80520646); instance of (P31);
1434 human (Q5)]
1435 5. [Abraham H. Albertson (Q80520646); archives at (P485);
1436 University of Washington Libraries (Q7896575)]

1437 Write a numbered list with the triples above, where each of
1438 the triples is followed by YES if the triple is represented in
1439 the text, and NO otherwise. Follow your 'YES' or 'NO' answer
1440 for each triple with a brief explanation.

1441 Below is the response generated by Meta-Llama-3.1-405B when prompted with the example
1442 above:
1443

1444 1. [Abraham H. Albertson (Q80520646); educated at (P69);
1445 Columbia Graduate School of Architecture, Planning and
1446 Preservation (Q5149701)] - **YES** - *The text explicitly states
1447 that Abraham H. Albertson graduated from the Columbia School
1448 of Architecture.*
1449 2. [Abraham H. Albertson (Q80520646); occupation (P106);
1450 architect (Q42973)] - **YES** - *Although the text does not directly
1451 state that Abraham H. Albertson was an architect, it can be
1452 inferred from the fact that he graduated from the Columbia
1453 School of Architecture.*
1454 3. [Abraham H. Albertson (Q80520646); country of
1455 citizenship (P27); United States of America (Q30)] - **NO** -
1456 *The text does not provide any information about Abraham H.
1457 Albertson's country of citizenship.*
1458 4. [Abraham H. Albertson (Q80520646); instance of (P31);

1458
1459 human(Q5)] - **YES** - *It can be inferred from the context that*
1460 *Abraham H. Albertson is a human being, as he is described as*
1461 *graduating from a school and having a war experience.*
1462 5. [Abraham H. Albertson(Q80520646); archives at(P485);
1463 University of Washington Libraries(Q7896575)] - **NO** - *The text*
1464 *does not provide any information about the archives of Abraham*
1465 *H. Albertson.*

1466 C.1.2 PROMPT FOR TRIPLE DEPRECATION
1467

1468 The following prompt is used to verify whether the textual passage supports the deprecation of a set
1469 of triples from the KG:

1470
1471 You are given the following text:
1472
1473 <TEXT>
1474
1475 Does this text imply that the following triples are no longer
1476 valid or do not hold at the moment the text was written?
1477
1478 <TRIPLES_LIST>
1479
1480 Write a numbered list with the triples above, where each of the
1481 triples is followed by YES if the text implies that the triple
1482 is no longer valid or do not hold, and NO otherwise. Follow
1483 your 'YES' or 'NO' answer for each of the triples with a brief
1484 explanation.

1485 The following is a concrete example of the *triple deprecation* prompt outlined above:
1486

1487 You are given the following text:
1488
1489 *Kate Lehrer (born Kate Tom Staples; December 17, 1937) is an*
1490 *American writer, novelist and book reviewer from Washington,*
1491 *D.C., and a panelist on the Diane Rehm Book Club on National*
1492 *Public Radio. She was married to fellow writer and journalist*
1493 *Jim Lehrer from 1960 until his death in 2020.*

1494 Does this text imply that the following triples are no longer
1495 valid or do not hold at the moment the text was written?
1496
1497 1. [Kate Lehrer(Q6375599); spouse(P26); Jim Lehrer(Q931148)]
1498 2. [Jim Lehrer(Q931148); spouse(P26); Kate Lehrer(Q6375599)]
1499 3. [Jim Lehrer(Q931148); residence(P551); Washington,
1500 D.C. (Q61)]

1501 Write a numbered list with the triples above, where each of the
1502 triples is followed by YES if the text implies that the triple
1503 is no longer valid or do not hold, and NO otherwise. Follow
1504 your 'YES' or 'NO' answer for each of the triples with a brief
1505 explanation.

1506
1507 Below is the response generated by Meta-Llama-3.1-405B when prompted with the example
1508 above:
1509

1510 1. [Kate Lehrer(Q6375599); spouse(P26); Jim Lehrer(Q931148)]
1511 - **YES** - *The text states that Kate Lehrer was married to Jim*

1512 Lehrer ‘until his death in 2020’, implying that the marriage
1513 ended with his death, and therefore the spouse relationship is
1514 no longer valid.
1515 2. [Jim Lehrer(Q931148); spouse(P26); Kate Lehrer(Q6375599)]
1516 - **YES** - *Similar to the previous triple, the text implies that*
1517 *the marriage between Jim Lehrer and Kate Lehrer ended with his*
1518 *death, so the spouse relationship is no longer valid.*
1519 3. [Jim Lehrer(Q931148); residence(P551); Washington,
1520 D.C. (Q61)] - **NO** - *The text does not provide any information*
1521 *about Jim Lehrer’s residence, let alone imply that he no longer*
1522 *resides in Washington, D.C. It only mentions that Kate Lehrer*
1523 *is from Washington, D.C.*

1524 1525 C.2 HUMAN ANNOTATION PROCESS

1526 To evaluate the accuracy of automatic LLM annotations, we employ two expert annotators (re-
1527 searchers holding a PhD with a background in computer science) to manually label a subsample
1528 of 500 (100 triples per each TKGU type defined in Section 3) triples from 166 randomly selected
1529 passages from our EMERGE dataset. To ensure consistency and avoid ambiguity, the annotators
1530 follow a detailed set of guidelines described in Section C.2.1.
1531

1532 C.2.1 ANNOTATION GUIDELINES

1533 *Annotators were provided with the following guideline:*

1534 For each annotation instance, you are provided with a textual passage, a triple, and an assess-
1535 ment type, which can be either *assert* or *deprecate*. For *assert* assessments, respond YES if the
1536 triple can be directly or indirectly inferred from the passage, and NO if it is not supported by
1537 the textual knowledge. For *deprecate* assessment, respond YES if the triple can be deprecated
1538 based on information present or implied in the passage, and NO otherwise. Take into account
1539 the following considerations when annotating for *assert* assessment type:

- 1540 1. The triple may not be factually correct at the time the text was written, but it ex-
1541 presses a fact that holds true at some other point in time. For example, the triple
1542 *⟨Barack Obama, president of, United States⟩* should be assessed YES for the text pas-
1543 sage “Barack Obama served as the 44th President of the United States from 2009 to
1544 2017”.
- 1545 2. Use common world knowledge and reasoning to induce triples from textual passage.
1546 For example, the triple *⟨Renault, headquarters in, France⟩* should be assessed YES for
1547 the text passage “The headquarters of Renault are located in Boulogne-Billancourt, a
1548 suburb of Paris.”, as Paris is located in France.
- 1549 3. Mark with NO any concrete fact that cannot be inferred from text, even if some of
1550 the entities appear in the passage. For example, the triple *⟨John Smith, participant*
1551 *in, Portland Climate Action Group protest⟩* should be assessed NO for the passage
1552 “Several members of the Portland Climate Action Group gathered downtown to protest
1553 against deforestation and climate inaction.”, as its factuality cannot be reliably inferred
1554 from the text.
- 1555 4. Assess with NO the triples that cannot be reliably inferred from a textual passage. For
1556 example, the triple *⟨David Bronkie, sibling, Eva Bronkie⟩* should be assessed as NO for
1557 the passage: “David Bronkie and Eva Bronkie co-founded a sustainable home goods
1558 business focused on eco-friendly candle kits.”, since the sibling relationship cannot be
1559 reliably inferred from the text (e.g., sharing the same last name).

1560 Take into account the following considerations when annotating for *deprecate* assessment type:

- 1561 1. The deprecation of a triple should be valid from the information provided in the pas-
1562 sage and not the current status of the knowledge. For example, the triple *⟨ Donald*
1563 *Trump, president of, United States ⟩* should be assessed with YES for the passage “Joe

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578

Biden is the President of the United States, having taken office recently and begun his tenure with notable public appearances and speeches.”, despite the fact that Donald Trump may be a current president of United States.

2. The depreciation of a triple might not be explicitly stated in the text, but can be implied. For example, the depreciation of the triple $\langle \text{Hans Rausing}, \text{spouse}, \text{Julia Rausing} \rangle$ should be assessed as YES for the passage “Julia Rausing, the philanthropist and business heiress, passed away on April 18, 2024, at the age of 63 after a long battle with cancer. She is survived by her husband, Hans Rausing, and their family.”, since the marital relationship is no longer current due to Julia Rausing’s death, which implies that the triple is deprecated.
3. Assess with NO any triples whose depreciation can not be reliably inferred from text, even if some of the entities appear in the text.

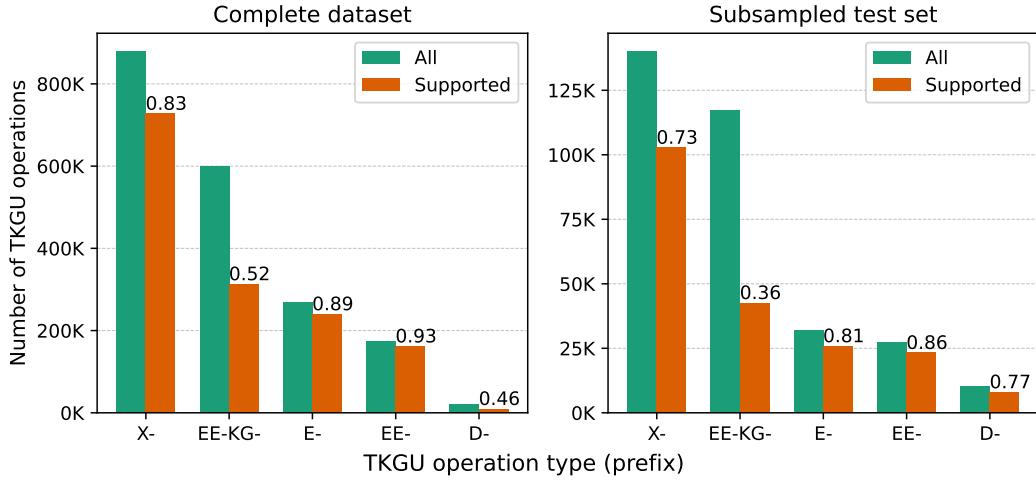


Figure 6: The ratio of TKGU operations supported by the LLM to the total number of TKGU operations mapped to textual passages during the alignment process.

C.2.2 ANNOTATION AGREEMENT

We report annotation agreement between the two human annotators (H - H Cohen’s κ), as well as between each human annotator and the LLM ($H1$ -LLM Cohen’s κ and $H2$ -LLM Cohen’s κ) in Table 5. The Cohen’s κ scores indicate strong agreement (0.6–0.8) to almost perfect agreement (> 0.8). In addition, we compute Fleiss’ κ ($H+LLM$ Fleiss’ κ) and Krippendorff’s α ($H+LLM$ Kripp. α) to assess agreement among all three annotators, both humans and the LLM. Consistent with Cohen’s κ , these metrics also show strong to almost perfect agreement. This supports the use of the evaluated Meta-Llama-3.1-405B LLM to annotate full dataset using the prompts described in the Appendix C.1.

C.3 TRIPLE ANNOTATION STATISTICS

Figure 6 illustrates the ratio of triples aligned with textual passages during the *alignment* step described in Section 4.2 that were marked by automatic LLM annotations – using the prompts detailed in Section C.1 – as not representative of the passages. This ratio is different between the *complete* and *subsampled* dataset used during testing. The reason is that during subsampling we retain instances with supported by LLM D-Triples operations (see Section 5). Additionally, we observe a lower fraction of EE-KG-Triples supported by the LLM. This occurs because EE-KG-Triples include all entities in the KG, many of which are unrelated to the passage content but are connected to emerging entities mentioned in the text. Consequently, these triples are inherently less likely

1620
1621 Table 5: Annotation agreement per TKGU operation and overall. Columns show pairwise Co-
1622 hen’s κ between humans (H-H) and between each human and the LLM (H1-LLM, H2-LLM), as
1623 well as multi-rater agreement including all three annotators (H+LLM) measured with Fleiss’ κ and
1624 Krippendorff’s α .

1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 TKGU Operation	1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 H-H Cohen’s κ	1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 H1-LLM Cohen’s κ	1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 H2-LLM Cohen’s κ	1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 H+LLM Fleiss’ κ	1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 H+LLM Kripp. α
X-Triples	0.718	0.649	0.637	0.668	0.669
E-Triples	0.750	0.698	0.750	0.732	0.733
EE-Triples	0.680	0.811	0.863	0.784	0.785
EE-KG-Triples	0.880	0.840	0.761	0.827	0.827
D-Triples	0.771	0.675	0.610	0.687	0.688
Overall	0.792	0.765	0.744	0.767	0.767

1634
1635
1636 Table 6: Statistics of our newly introduced EMERGE dataset, organized by KG snapshots (rows).
1637 For each snapshot, we report the number of *instances* and TKGU *operations* in both the *complete*
1638 *dataset* and the *subsampled test set*. The *KG statistics* section summarizes the number of entities,
1639 relation types, and triples in each KG snapshot.

1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Snapshot	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Complete dataset	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Subsampled test set	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 KG statistics		
	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Instances	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Operations	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Entities	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Rel. Types	1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 Triples
2019	37K	202K	5K	24K	5.96M
2020	31K	199K	5K	26K	6.14M
2021	40K	292K	5K	36K	6.34M
2022	30K	188K	5K	27K	6.54M
2023	26K	151K	5K	26K	6.67M
2024	32K	200K	5K	29K	6.80M
2025	33K	217K	5K	31K	6.93M

1653 to be supported by the passages. A promising future direction is to develop information extraction
1654 methods that rely not only on textual evidence to extract triples but also integrate this content
1655 with existing knowledge and patterns in the KG. Such an approach could be particularly beneficial
1656 for incorporating emerging entities in EE-KG-Triples, even when they are not supported by textual
1657 passages, into the broader KG.

1658

1659

1660 **D DATASET STATISTICS**

1661

1662 In this section we will present additional statistics of EMERGE.

1663

1664

1665 **D.1 OVERALL STATISTICS OF EMERGE**

1666

1667 Table 6 presents key statistics of our newly introduced EMERGE dataset, broken down by KG reference
1668 snapshots. For each snapshot, we report the number of *instances* and TKGU *operations* in both the full dataset and the subsampled test set. The table also summarizes *KG* snapshots statistics,
1669 including the number of entities, relation types, and triples in each snapshot. We observe that the
1670 number of entities, relation types, and triples increases over time, reflecting the growth of Wikidata
1671 and the addition of new relations to the KG schema. This evolving structure creates a challenging
1672 scenario for future models, which must recognize these changes in the KG and adapt their predictions
1673 accordingly.

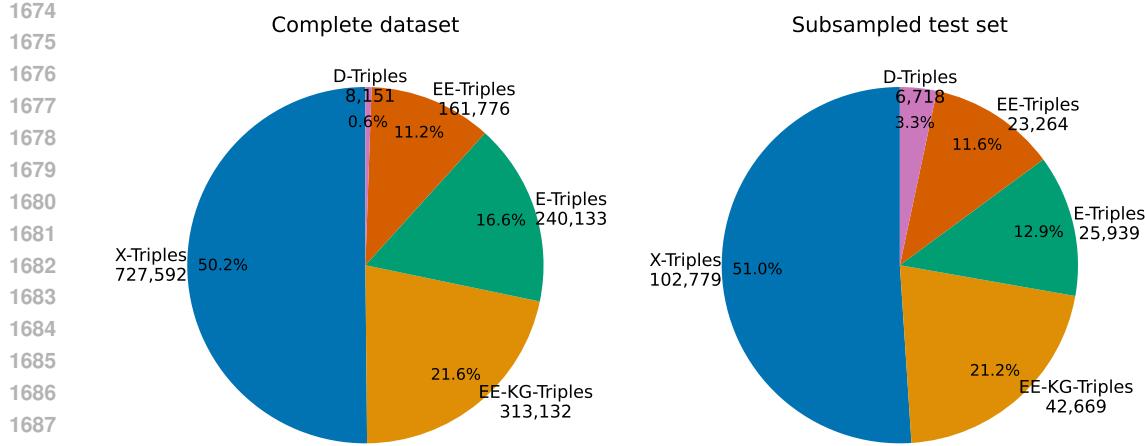


Figure 7: Distribution of TKGU operations defined in Section 3 in EMERGE. The left subgraph shows the full dataset, while the right one shows the subsampled test set (see Section 5). In the test set, D-Triples are retained at higher frequency to ensure sufficient evaluation, while other TKGU operation types reflect the original dataset distribution.

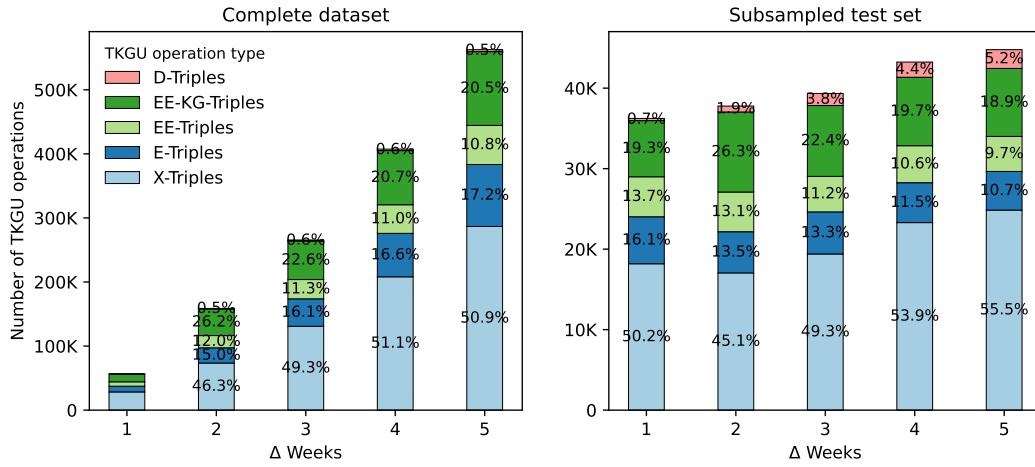


Figure 8: Distribution of TKGU operations across KG deltas up to 5 weeks defined in EMERGE.

D.2 NUMBER OF TKGU OPERATIONS AND THEIR DISTRIBUTION

Figure 7 illustrates the distribution of KG update operations for each TKGU type defined in Section 3. We report on both the complete dataset (left subgraph) and the subsampled test set (right subgraph). Furthermore, we display both the number as well as the percentage the operations of each of the TKGU types represent in EMERGE. This distribution is very similar between the complete dataset and subsampled test set, except for D-Triples, which were retained at higher frequency in the test set to ensure sufficient evaluation (see Section 5). Additionally, Figure 8 shows the distribution of TKGU update operations across temporally increasing weekly KG deltas. In the *Complete dataset* (left subplot), the number of TKGU operations increases with larger deltas. A similar trend is visible in the *subsampled test set* (right subplot), although the growth is less pronounced. This is due to our subsampling procedure, which retains only 1,000 instances per delta (see Section 5), resulting in a more uniform distribution of operations across deltas.

1728
1729

E QUALITATIVE ANALYSIS

1730 In this section, in Tables 7–12 we present the five frequent factual triples from EMERGE for each
1731 of the TKGU operation types, with an example of corresponding textual passage. The goal is to
1732 highlight representative cases that illustrate both the contents of the benchmark and the challenges
1733 it poses. The information in the tables contains the KG snapshot (*Snap.*) used to compute weekly
1734 knowledge deltas aligned with each passage. We also report the number of occurrences of the triple
1735 in the *Triple* column within EMERGE (#), along with an example passage. The emerging entities
1736 in TKGU operations appear in bold. Due to space constraints, we selected the shortest passages;
1737 however, in EMERGE, passages consist of full Wikipedia paragraphs.

1738 Our main observation is that the derived TKGU operations are closely aligned with the primary
1739 events occurring immediately after each KG snapshot (all snapshots are taken on January 1st of the
1740 corresponding year). We also note that the resulting triples are highly specific to the Wikidata KG
1741 structure. This is particularly evident in Table 10, which shows examples of EE-KG-Triples, where
1742 an emerging entity must be connected to the existing KG. Consequently, we believe a promising
1743 future direction is to develop information extraction models that consider KG structure when proposing
1744 knowledge updates in it.

1745 Additionally, to illustrate the effectiveness of using an LLM (Meta-Llama-3.1-405B) to verify
1746 that all TKGU operations can be derived from their corresponding textual passages during the
1747 *curation* step described in Section 4.2, we present the most frequent factual triples from the EE-KG-
1748 Triples TKGU operation in EMERGE that were marked as *not supported* by the LLM in Table 11.
1749 These examples highlight triples that occurred frequently but were flagged because the LLM deter-
1750 mined that their source textual passages did not support them. None of these triples are grounded in
1751 the corresponding text, demonstrating the reliability of the LLM-based validation process.

1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

1782

1783 Table 7: Example entries of the most frequent X-Triples TKGU operation instances in EMERGE,
1784 showing the snapshot (Snap.), triple, and number of instances (#).

1785

1786

Snap.	#	Triple	Example Passage
2021	834	⟨Donald Trump; candidacy in election; 2020 United States presidential election⟩	Over the span of the 2020 presidential election, RSBN’s coverage of Donald Trump’s campaign rallies grossed over 127 million views on YouTube.
2021	827	⟨2020 United States presidential election; candidate; Donald Trump⟩	In 2020, Pletts voiced support for Donald Trump and the Republican Party in the 2020 United States presidential election and Senate elections.
2021	671	⟨Joe Biden; candidacy in election; 2020 United States presidential election⟩	In September 2020, Kennedy Kent endorsed Republican President Donald Trump for re-election over Democratic nominee Joe Biden.
2021	666	⟨2020 United States presidential election; candidate; Joe Biden⟩	Despite being divorced, she remains good friends with her ex-husband, and she supported Joe Biden and Kamala Harris in the 2020 election.
2021	586	⟨midfielder; sport; association football⟩	“Niko Rak” (born 26 July 2003) is a Croatian footballer who plays for Šibenik as a midfielder.

1805

1806

1807

1808

1809

1810

1811

1812

1813

1814

Table 8: Example entries of the most frequent E-Triples TKGU operation instances in EMERGE, showing the snapshot (Snap.), triple, and number of instances (#).

1815

1816

Snap.	#	Triple	Example Passage
2021	315	⟨Joe Biden; position held; President of the United States⟩	On 20 January 2021, Joe Biden was sworn in as 46th President of the United States.
2023	204	⟨Kevin McCarthy; position held; Speaker of the United States House of Representatives⟩	On January 3, 2023, at the beginning of the 118th Congress, Boebert voted for Jim Jordan to be the U.S. House Speaker, in rebuke of House Minority Leader Kevin McCarthy.
2020	168	⟨Abu Mahdi al-Muhandis; military branch; Popular Mobilization Forces⟩	Abu Mahdi al-Muhandis returned to Iraq following the withdrawal of US troops (December 2011) to head the Kata’ib Hezbollah militia.; he then became deputy chief of the Popular Mobilization Forces.
2024	164	⟨Houthi movement; country; Yemen⟩	On 28 March 2021, the Houthis forced 13 Jews to leave Yemen, they only allowed four elderly Jews to live in Yemen.
2020	138	⟨Qasem Soleimani; place of death; Baghdad⟩	Soleimani was assassinated in a targeted U.S. drone strike on 3 January 2020 in Baghdad, which was approved by President Donald Trump on the grounds that Soleimani posed an “imminent threat” to American lives.

1835

1836

1837 Table 9: Example entries of the most frequent EE-Triples TKGU operation instances in
1838 EMERGE (emerging entities in bold), showing the snapshot (Snap.), triple, and number of instances
1839 (#).

Snap.	#	Triple	Example Passage
2021	848	⟨January 6 United States Capitol attack; significant person; Donald Trump⟩	She called for the impeachment of President Donald Trump, in wake of the 2021 storming of the United States Capitol.
2020	670	⟨Qasem Soleimani; significant event; assassination of Qasem Soleimani⟩	He was killed by a targeted U.S. drone strike at the Baghdad International Airport on 3 January 2020, which also killed Iranian Armed Forces Major General Qasem Soleimani.
2022	317	⟨Dawn FM; performer; The Weeknd⟩	In 2022 the group also received credit for co producing songs off The Weeknd's fifth studio album Dawn FM.
2025	291	⟨2025 New Orleans truck attack; located in the administrative territorial entity; New Orleans⟩	2025 New Orleans truck attack: President Joe Biden has been briefed on the attack and has been in touch with New Orleans Mayor to offer support.
2023	72	⟨Flowers; performer; Miley Cyrus⟩	The chart's current number one as of the issue dated January 28, 2023, is "Flowers" by Miley Cyrus

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869 Table 10: Example entries of the most frequent EE-KG-Triples TKGU operation instances in
1870 EMERGE (emerging entities in bold), showing the snapshot (Snap.), triple, and number of instances
1871 (#).

Snap.	#	Triple	Example Passage
2021	3149	⟨January 6 United States Capitol attack; located in the administrative territorial entity; Washington, D.C.⟩	January 6 United States Capitol attack: The Proud Boys posted messages boasting and taking credit for causing "absolute terror".
2020	1097	⟨assassination of Qasem Soleimani; instance of; assassination⟩	Assassination of Qasem Soleimani: the president called for restraint and said the events in Iraq were the result of previous "terrorist acts".
2025	991	⟨Golden Age of Argentine cinema; part of; history of film⟩	"Volver a vivir" is a 1941 Argentine film of the Golden Age of Argentine cinema.
2024	282	⟨South Africa v. Israel; charge; genocide⟩	In 2023-24, he was appointed as a member of the South African legal team arguing "South Africa v. Israel" regarding the Genocide Convention.
2019	179	⟨All Elite Wrestling; instance of; business⟩	On January 1, 2019 Cody Rhodes unveiled a new promotion; All Elite Wrestling, in which he, along with Matt and Nick Jackson, will serve as Executive Vice President.

1890
1891
1892
1893
1894
1895

Table 11: Representative examples of the most frequent EE-KG-Triples TKGU operation instances in **EMERGE filtered out by the Meta-Llama-3.1-405B curator**. Emerging entities are in bold. Each row shows the snapshot (Snap.), triple, and number of occurrences (#). None of these triples are supported by the corresponding textual passages, illustrating the effectiveness of the LLM-based filtering.

Snap.	#	Triple	Example Passage
2021	1174	⟨Proud Boys; significant event; January 6 United States Capitol attack ⟩	Trump supporters infiltrated Capitol Hill in Washington DC., 5 people killed.
2022	158	⟨ Dawn FM ; distribution format; LP record⟩	In 2022 the group also received credit for co producing songs off The Weeknds fifth studio album Dawn FM.
2024	123	⟨ 2024 Haneda Airport runway collision ; destination point; Niigata Airport⟩	2024 Haneda Airport runway collision: All flights in and out of Haneda were suspended following the accident; operations currently remain suspended.
2019	55	⟨ All Elite Wrestling ; legal form; privately held company⟩	On January 1, 2019 Cody Rhodes unveiled a new promotion; All Elite Wrestling, in which he, along with Matt and Nick Jackson, will serve as Executive Vice President.
2019	29	⟨ @world_record_egg ; country; United Kingdom⟩	@world_record_egg is an account on social media platform Instagram, notable for holding the world records for both the most-liked Instagram post and most liked online post on any media platform in history.

1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

1944
1945
1946
1947

Table 12: Example entries of the most frequent D-Triples TKGU operation instances in EMERGE (emerging entities in bold), showing the snapshot (Snap.), triple, and number of instances (#).

Snap.	#	Triple	Example Passage
2024	88	⟨Adam Peters ; member of sports team; San Francisco 49ers⟩	Peters joined the Denver Broncos as a scout in 2009. He was promoted to assistant director of college scouting in July 2014 and to director of college scouting in 2016. He was a member of the team that won Super Bowl 50 in 2015.
2021	79	⟨Parler ; distributed by; Google Play⟩	After complaints that Parler was used to coordinate the 2021 storming of the U.S. Capitol, Apple and Google removed Parler’s mobile app from their app stores. Parler went offline on January 10, 2021 at 11:59 PM (PST) after Amazon Web Services canceled its hosting services.
2021	75	⟨Mike Pence ; position; Vice President of the United States⟩	“Marlon Bundo”, also known as “Bunny of the United States” (“BOTUS”), is a rabbit, belonging to the family of Mike Pence, the 48th and former Vice President of the United States.
2025	63	⟨Vice President of the United States ; position holder; Kamala Harris⟩	West is the brother-in-law of former Vice President Kamala Harris. He served as an advisor to her 2024 presidential campaign.
2020	43	⟨European Union ; has part(s); United Kingdom⟩	Chris Davies was the chairman (2019 – 2020) - until the United Kingdom left the European Union.

F ADDITIONAL EXPERIMENTAL RESULTS

F.1 PRECISION AND F1 SCORES

In this section, we present precision and F1 results for both *closed information extraction* (ReLiK cIE; Tables 13–14) and *relation extraction* (ReLiK RE and EDC+; Tables 15–16). Closed information extraction performance is measured via exact triple matching against the ground truth. Relation extraction is evaluated using the completeness-score approximation with a threshold $\phi = 0.9$ (see Section B.2 for the formal definition).

Table 13: *Precision* for the closed IE model ReLiK cIE (i.e., the extracted triples are linked to the KG) across KG snapshots, evaluated using the TKGU operations defined in Section 3.

TKGU	Model	2019	2020	2021	2022	2023	2024	2025
X-Triples	ReLiK cIE	45.8	31.4	35.1	37.2	33.3	34.0	31.2
E-Triples	ReLiK cIE	2.6	3.2	1.9	2.9	3.3	2.7	2.7

Table 14: *F1 score* for the closed IE model ReLiK cIE (i.e., the extracted triples are linked to the KG) across KG snapshots, evaluated using the TKGU operations defined in Section 3.

TKGU	Model	2019	2020	2021	2022	2023	2024	2025
X-Triples	ReLiK cIE	25.9	21.5	21.6	20.5	21.2	18.6	20.1
E-Triples	ReLiK cIE	4.5	5.4	3.3	4.8	5.5	4.4	4.6

1998
1999
2000

Table 15: *Precision* (measured using the completeness score) for IE models that do not link extracted triples to the KG, evaluated across KG snapshots on the TKGU operations defined in Section 3.

TKGU	Model	2019	2020	2021	2022	2023	2024	2025
X-Triples	EDC+ Mistral-7b	4.7	3.2	4.9	4.1	5.2	3.7	4.9
	EDC+ Gemma-7b	3.2	2.8	3.1	2.5	3.2	2.5	3.4
	ReLiK RE	9.2	8.2	8.1	9.2	8.6	9.2	9.9
E-Triples	EDC+ Mistral-7b	2.5	2.3	1.9	2.3	2.4	2.3	2.1
	EDC+ Gemma-7b	2.0	1.5	1.3	1.6	1.8	1.5	1.5
	ReLiK RE	2.4	2.0	1.9	1.9	1.9	1.7	1.5
EE-Triples	EDC+ Mistral-7b	1.9	2.3	1.1	1.8	2.1	1.6	1.8
	EDC+ Gemma-7b	1.4	1.5	0.9	1.4	1.7	1.3	1.4
	ReLiK RE	1.7	1.9	1.0	2.4	2.2	1.7	1.6
EE-KG-Triples	EDC+ Mistral-7b	1.6	2.1	1.8	1.8	1.9	1.5	1.8
	EDC+ Gemma-7b	2.2	2.7	1.6	1.9	2.4	2.0	2.0
	ReLiK RE	0.3	0.7	0.8	0.5	0.6	0.6	0.7
D-Triples	EDC+ Mistral-7b	8.5	9.8	14.4	9.6	15.5	5.8	11.6
	EDC+ Gemma-7b	5.0	6.9	8.3	8.9	8.6	3.3	6.4

2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018

Table 16: *F1 score* (measured using the completeness score) for IE models that do not link extracted triples to the KG, evaluated across KG snapshots on the TKGU operations defined in Section 3.

TKGU	Model	2019	2020	2021	2022	2023	2024	2025
X-Triples	EDC+ Mistral-7b	6.3	4.5	6.7	5.4	7.2	5.0	6.3
	EDC+ Gemma-7b	4.5	4.2	4.4	3.5	4.5	3.5	4.5
	ReLiK RE	13.5	12.3	12.1	12.6	12.4	12.4	13.3
E-Triples	EDC+ Mistral-7b	4.5	4.0	3.3	4.0	4.3	4.2	3.8
	EDC+ Gemma-7b	3.5	2.8	2.3	2.8	3.2	2.8	2.7
	ReLiK RE	4.3	3.7	3.4	3.4	3.4	3.0	2.8
EE-Triples	EDC+ Mistral-7b	3.4	4.0	2.0	3.2	3.7	2.8	3.2
	EDC+ Gemma-7b	2.6	2.7	1.6	2.6	3.1	2.4	2.6
	ReLiK RE	3.1	3.5	1.8	4.4	4.0	3.1	2.9
EE-KG-Triples	EDC+ Mistral-7b	3.0	3.7	2.9	3.3	3.5	2.8	3.4
	EDC+ Gemma-7b	3.7	4.0	1.8	3.1	3.8	3.1	3.0
	ReLiK RE	0.5	1.3	1.3	0.8	1.0	1.1	1.2
D-Triples	EDC+ Mistral-7b	7.8	9.8	10.0	8.5	13.9	4.7	9.9
	EDC+ Gemma-7b	5.2	8.4	8.3	9.6	9.6	4.2	6.6

2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035

G QUALITATIVE ANALYSIS OF THE RESULTS

Table 17 presents three instances from EMERGE, together with the predictions of the evaluated models described in Section 5.1. We observe that while most model predictions are semantically correct, they often do not correspond to the TKGU operations that capture actual changes in the KG. For instance, in *passage 1*, the predicted *D-Triples* are reasonable but differ from the triple actually deprecated in the KG: `<CLC; has part; Elkie>`. This discrepancy is expected, as current state-of-the-art information extraction models are largely unaware of the structure and content of KGs (see Table 1). We observe similar mismatches for the X/E/EE/EE-KG-Triples operations in *passage 2* and *passage 3*. Moreover, the LLM-driven EDC+ method exhibits a high degree of hallucination, particularly for **EE-KG-Triples**, where an emerging entity must be connected to the KG. Since EDC+ does not interact with the KG, its predicted emerging entities often fail to correspond to the actual ones. For instance, in *passage 3*, the true emerging entity is Puttsburg Maulers, yet both *EDC+ Mistral-7b* and *EDC+ Gemma-7b* instead generate EE-KG-Triples around Kirby Wilson. In addition, many of the generated triples are not supported by the passage itself. We hypothesize

2052 that enabling these models to interact with KG content and structure could mitigate such errors, as
2053 predictions would then be grounded not only in the text but also aligned with existing KG contents.
2054

2055

2056

2057

2058

2059

2060

2061

2062

2063

2064

2065

2066

2067

2068

2069

2070

2071

2072

2073

2074

2075

2076

2077

2078

2079

2080

2081

2082

2083

2084

2085

2086

2087

2088

2089

2090

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

2101

2102

2103

2104

2105

2106
2107 Table 17: Table showcasing three instances (passages aligned to Ground truth TKGU operations) of
2108 EMERGE. We also display the predictions of the models used in our experiments (see Section 5.1).
2109 While these predictions are often correct, they frequently fail to align with the existing knowledge
2110 and structure of the KG.

2111	Passage 1	On February 3, 2021, Cube Entertainment confirmed Elkie’s departure from CLC, and her contract with the company has been terminated.
2112	Ground truth	D-Triples: ⟨CLC; has part; Elkie⟩
2113	EDC+ Mistral-7b	D-Triples: ⟨Elkie; member of musical group; CLC⟩
2114	EDC+ Gemma-7b	D-Triples: ⟨CLC; employee; Elkie⟩
2115	ReLiK RE	–
2116	ReLiK cIE	–
2117	Passage 2	“Cancilla liliiformis” is a species of sea snail, a marine gastropod mollusk, in the family Mitridae, the miters or miter snails.
2118	Ground truth	EE-Triples: ⟨Cancilla liliiformis; taxon rank; species⟩ X-Triples: ⟨Mitridae; taxon rank; family⟩, ⟨Gastropoda; parent taxon; Mollusca⟩ EE-KG-Triples: ⟨Cancilla liliiformis; instance of; taxon⟩
2119	EDC+ Mistral-7b	X/E/EE-Triples: ⟨Cancilla liliiformis; species of; Mitridae⟩ EE-KG-Triples: ⟨Cancilla liliiformis; instance of; marine organism⟩, ⟨Cancilla liliiformis; instance of; mollusk⟩
2120	EDC+ Gemma-7b	X/E/EE-Triples: ⟨Cancilla liliiformis; species of; Mitridae⟩ ⟨Cancilla liliiformis, instance of, Gastropoda⟩ EE-KG-Triples: ⟨Cancilla liliiformis, country of citizenship, France⟩, ⟨Cancilla liliiformis, place of birth, France⟩, ⟨Cancilla liliiformis, instance of, Human⟩
2121	ReLiK RE	X/E/EE/EE-KG-Triples: ⟨Cancillopsis liliiformis; subclass of; Mitridae⟩, ⟨Cancilla liliiformis; subclass of; Mitridae⟩
2122	ReLiK cIE	X/E/EE/EE-KG-Triples: ⟨Sea snail; subclass of; gastropod⟩, ⟨Mitridae; subclass of; gastropod⟩
2123	Passage 3	Pittsburgh Maulers (USFL): On January 20, 2022, Wilson was named Head Coach and General Manager of the Pittsburgh Maulers of the United States Football League (USFL).
2124	Ground truth	EE-Triples: ⟨Pittsburgh Maulers; league; United States Football League⟩ EE-KG-Triples: ⟨Pittsburgh Maulers; country; United States of America⟩, ⟨Pittsburgh Maulers; sport; American football⟩
2125	EDC+ Mistral-7b	X/E/EE-Triples: ⟨Pittsburgh Maulers, head coach, Kirby Wilson⟩ EE-KG-Triples: ⟨Kirby Wilson; educated at; University of Pittsburgh⟩, ⟨Kirby Wilson; place of birth, Pittsburgh⟩, ⟨Kirby Wilson; country of citizenship; United States⟩, ⟨Kirby Wilson; instance of; human⟩
2126	EDC+ Gemma-7b	X/E/EE-Triples: ⟨Pittsburgh Maulers; head coach; Kirby Wilson⟩, ⟨Pittsburgh Maulers; country; United States⟩ EE-KG-Triples: ⟨Kirby Wilson; educated at; University of North Texas⟩, ⟨Kirby Wilson; place of birth; Dallas⟩, ⟨Kirby Wilson; instance of; human⟩
2127	ReLiK RE	X/E/EE/EE-KG-Triples: ⟨Kirby Wilson; member of sports team; Pittsburgh Maulers⟩, ⟨Kirby Wilson; member of; Pittsburgh Maulers⟩, ⟨Kirby Wilson; member of sports team; Pittsburgh Maulers⟩
2128	ReLiK cIE	⟨Kirby Wilson; member of sports team; Pittsburgh Maulers⟩, ⟨Wilson; member of sports team; Pittsburgh Maulers⟩, ⟨Head Coach; member of sports team; Pittsburgh Maulers⟩, ⟨Pittsburgh Maulers; organizer; USFL⟩

2160 **H EDC+ EXECUTION**
2161

2162 **H.1 EDC+ EXECUTION TIME**
2163

2164 To generate predictions on the subsampled test set (see above), we run EDC+ with the
2165 `Mistral-7B-Instruct-v0.2` and `gemma-7b` LLMs on two H100 GPUs for 24 hours.
2166

2167 **H.2 EDC+ PROMPTS**
2168

2169 The following prompt is designed to identify all the operations to update the KG defined in Section 3.
2170 Concretely, it allows to identify triples explicitly mentioned in text under *Triples in text* category.
2171 This includes *X-Triples*, *E-Triples*, and *EE-Triples*. It also allows to classify *Triples in text* in those
2172 that should be added to the KG (i.e., with the ADD tag), and those that should be deprecated (i.e.,
2173 with the DEPRECATE tag). This way, the prompt also facilitates the identification of KG triples
2174 that may need to be deprecated (i.e., *D-Triples*). Finally, the prompt allows to detect *EE-KG-Triples*
2175 under *Triples not in text* category, by asking LLM to identify triples with only one single entity (head
2176 or tail) mentioned in text, and the other entity existing in the KG.
2177

2178 Your task is to transform the given text into a semantic graph
2179 in the form of a list of triples. Two sets of triples are
2180 to be extracted: 'Triples in text', which contain triples
2181 relating entities mentioned in text in the form of [Entity1,
2182 Relationship, Entity2, Action], where action indicates if the
2183 triple has to be added (action 'ADD') or deprecated (action
2184 'DEPRECATE') from the graph based on the knowledge in text.
2185 The second set of triples is called 'Triples not in text', and
2186 consists of triples with one entity (head or tail) mentioned in
2187 text and the other entity not mentioned in text but potentially
2188 existing in the graph.
2189 In your answer, please strictly only include the triples and do
2190 not include any explanation or apologies.
2191 Here are some examples:
2192

2193 <FEW_SHOT_EXAMPLES>

2194 Now please extract triplets from the following text.
2195

2196 Text: <INPUT_TEXT>

2197 **I RELiK EXPERIMENTAL CONFIGURATION**
2198

2199 To generate predictions, we run ReLiK on each KG snapshot independently. In each run,
2200 ReLiK is provided with the dictionary of entities and relations specific to that snapshot.
2201 For relation encoding, we use the pre-trained ReLiK model available on Hugging Face:
2202 `relik-ie/encoder-e5-small-v2-wikipedia-relations`. These relation encodings are used by both ReLiK RE and ReLiK cIE. For each snapshot, we also encode the corresponding KG entities using the model `relik-ie/encoder-e5-small-v2-wikipedia-matryoshka`.
2205

2206 For prediction, we use the pre-trained `relik-ie/relik-relation-extraction-large` model for ReLiK RE, and the pre-trained `relik-ie/relik-cie-large` model for ReLiK cIE.
2207

2208 Running ReLiK on the subsampled EMERGE test set takes about 5 hours on a single A100 GPU.
2209

2210 **J WIKIDATA QUALIFIERS TO DETECT DEPRECATION OF TRIPLES**
2211

2212 The following is the list of Wikidata qualifiers we use to detect the deprecation of triples when
2213 creating EMERGE:

2214 1. P582: end time.
2215 2. P1326: latest date.
2216 3. P576: dissolved, abolished or demolished date.
2217 4. P570: date of death.
2218 5. P730: service retirement.
2219 6. P2032: work period (end).
2220 7. P2669: discontinued date.
2221 8. P3999: date of official closure.
2222 9. P7125: date of the latest one.

2223 **K LIMITATIONS AND FUTURE WORK**

2224 In this work, we focus specifically on changes to the KG that reflect the introduction or modification
2225 of factual knowledge. We do not account for structural or curation-related changes that a KG may
2226 undergo, such as schema adjustments, property reorganization, or entity merging. These types of
2227 changes are often independent of new information appearing in external sources like Wikipedia and
2228 are typically driven by internal quality control or ontology refinement processes. While important
2229 for maintaining the integrity and usability of the KG, such changes fall outside the scope of our
2230 current study.

2231 In this work, we focus on leveraging external textual sources to enhance KGs. However, textual data
2232 represents only one type of external knowledge. Other modalities—such as video (e.g., podcasts),
2233 images, and audio—also contain rich, complementary information that can contribute to KG enrichment.
2234 As such, a promising direction for future research is to explore the integration of knowledge
2235 from these multimodal sources to address this limitation.

2236 During the generation of EMERGE, we use the same temporal delta window for both, the extraction
2237 of changes in Wikidata and the emerging passages from Wikipedia. However, certain pieces of
2238 knowledge do not always appear within the same time frame in the two sources. For example,
2239 events such as Brexit or the election of a president are often documented in Wikipedia months or
2240 even years before they are incorporated into the Wikidata knowledge graph. In future work, we plan
2241 to investigate this temporal discrepancy between the two knowledge sources more thoroughly.

2242 Furthermore, this study restricts attention to triples in which both the subject and object are entities
2243 present in the entity catalog. Nonetheless, numerous valuable relations involve literals as objects,
2244 such as dates of birth, lengths, sizes, or employee counts (Mesquita et al., 2019), which are not
2245 considered in the current work.

2246 Finally, a limitation of EMERGE is that it covers only the subset of Wikidata changes that can be
2247 reliably grounded in Wikipedia text. This stems from the fact that Wikidata is crowdsourced and not
2248 fully determined by Wikipedia content, meaning that many Wikidata updates have no corresponding
2249 textual evidence. In addition, EMERGE is restricted to Wikipedia paragraphs in which annotated
2250 entity mentions can be reliably identified through hyperlinks, as described in Section 4.1. As future
2251 work, an alternative dataset could be constructed using text-to-data generation methods (Hu et al.,
2252 2025; Edge et al., 2024; Hofer et al., 2024) to create a synthetic KG that mirrors all knowledge found
2253 in text, thereby achieving complete coverage of updates. While such an approach would ensure full
2254 alignment between text and KG, it would also introduce challenges such as potential errors in entity
2255 disambiguation and the substantial computational cost of generating an entire KG from text.

2256 Another direction for future work is to incorporate explicit start and end dates for TKGU operations
2257 that imply changes in the KG, such as triple addition and deprecation. In the current version of
2258 EMERGE, deprecated facts are identified through the delta interval, as our work primarily focuses
2259 on updating the KG at a specific point in time rather than modeling full temporal validity. Adding
2260 explicit temporal qualifiers would more precisely capture when a fact begins and ceases to hold,
2261 aligning EMERGE more closely with the way temporal information is handled in Wikidata. This
2262 extension would also enable richer modeling of fact evolution and support downstream methods that
2263 rely on explicit temporal boundaries.

2268 A further promising future direction is to develop information extraction methods that rely not only
2269 on textual evidence to extract triples, but also integrate this content with existing knowledge and
2270 patterns in the KG. Such an approach could be particularly beneficial for incorporating emerging
2271 entities in EE-KG-Triples, even when they are not supported by textual passages, into the broader
2272 KG.
2273

2274 L DATASET DOCUMENTATION: DATASHEET 2275

2276 We describe our dataset following the datasheets for datasets guidelines introduced in (Gebru et al.,
2277 2021), detailing its motivation, composition, collection process, and recommended uses. This doc-
2278 umentation supports transparency, reproducibility, and responsible dataset use in machine learning
2279 research.
2280

2281 L.1 MOTIVATION 2282

2283 **For what purpose was the dataset created?** The EMERGE dataset was created to address the
2284 lack of integration between changes in textual knowledge and their effect on knowledge graph con-
2285 tent. The proposed benchmark enables evaluation of KG updates driven by newly emerging knowl-
2286 edge in textual sources over temporally increasing KG deltas. Moreover, because the dataset is
2287 generated via an automatic annotation pipeline, it can be continuously extended to include more
2288 recent knowledge, thereby allowing evaluation of model robustness to ever-evolving and novel in-
2289 formation and KG structures. This contrasts with existing benchmarks (see Table 4 in the Appendix
2290 A.1), which are static in nature and unable to emulate the evolution of knowledge in textual and KG
2291 sources (columns *Evolution-KG* and *Evolution-Text* in Table 4). Furthermore, existing benchmarks
2292 do not cover all the necessary text-driven knowledge graph update (TKGU) operations necessary
2293 to keep them updated (columns *X-Triples*, *E-Triples*, *EE-Triples*, *EE-KG-Triples* and *D-Triples* in
2294 Table 4).
2295

2296 We expect EMERGE will encourage the research on methods that are not limited to extracting
2297 knowledge from textual sources, but also capable of effectively maintaining KGs by integrating that
2298 knowledge into existing KGs. This contrasts with current state-of-the-art IE methods (see Section 2
2299 and Table 1) limited to the extraction of knowledge purely from text without the ability to effectively
2300 integrate that knowledge into existing knowledge in KGs.
2301

2302 **Who created the dataset (e.g., which team, research group) and on behalf of which entity**
2303 **(e.g., company, institution, organization)?** The dataset was developed by academic researchers
2304 through an international, cross-institutional collaboration. The contributing researchers bring exten-
2305 sive expertise in information extraction methods and dataset construction.
2306

2307 **Who funded the creation of the dataset?** The dataset was created with funding from, among
2308 others, the highly prestigious European Union Marie Curie Actions Postdoctoral Grant.
2309

2310 L.2 COMPOSITION 2311

2312 **What do the instances that comprise the dataset represent (e.g., documents, photos, peo-
2313 ple, countries)?** The instances that comprise the dataset represent general-domain passages from
2314 Wikipedia, KG triples representing the knowledge contained in those passages, and TKGU opera-
2315 tions (see Section 3) with respect to the respective general-domain Wikidata KG snapshot.
2316

2317 **How many instances are there in total (of each type, if appropriate)?** Our EMERGE contains
2318 in total 233K instances, with a total of 1.4M TKGU operations: 727K X-Triples, 240K E-Triples,
2319 161K EE-Triples, 313K EE-KG-Triples, and 8K D-Triples.
2320

2321 **Does the dataset contain all possible instances or is it a sample (not necessarily random) of
2322 instances from a larger set?** We include a set with all possible instances that can be used for
2323 training. For testing (on which we report our results), we subsampled 1,000 instances per delta per
2324 snapshot.
2325

2322 **What data does each instance consist of?** Each of the instances in the dataset consists of a
2323 textual passage with an annotated set of entity mentions linked to a particular KG snapshot. In
2324 addition, each instance includes a list of triples together with the corresponding TKGU operations
2325 that update the KG snapshot, as described in Section 3. Each triple is further annotated with an
2326 LLM-based assessment indicating whether the knowledge it represents can be inferred from the
2327 textual passage. See Appendix C.1 for details on the prompt and examples. The dataset spans seven
2328 yearly KG snapshots covering 2019–2025. For each snapshot, TKGU updates are annotated over five
2329 progressively larger weekly KG deltas, thereby capturing different levels of knowledge staleness in
2330 the KG.

2331 **Is there a label or target associated with each instance?** Yes, the target consists of all the triples
2332 with corresponding TKGU operations associated with the textual passage of an instance. These
2333 operations specify the updates to be applied to a KG snapshot to ensure consistency with the textual
2334 passage.

2335 **Is any information missing from individual instances?** All the instances are consistently anno-
2336 tated. However, the triples involved in TKGU operations associated with a passage are restricted
2337 to the entities of mentions explicitly annotated with hyperlinks in Wikipedia (see Section 4.1 for
2338 further details on annotation process). As such, there might be TKGU operations not covered by our
2339 dataset. This is also discussed in the limitations sections (see Section K).

2340 **Are relationships between individual instances made explicit (e.g., users’ movie ratings, social
2341 network links)?** Yes, all the detected TKGU operations during the annotation process are made
2342 explicit. We further mark each of these operations as supported or no by the content of textual
2343 passage using LLM automatic annotation process described in Section 4.2.

2344 **Are there recommended data splits (e.g., training, development/validation, testing)?** Yes. We
2345 recommend training and validating models on earlier snapshots (e.g., from 2019 and 2020) and
2346 testing on later snapshots (i.e., from 2021–2025). This setup prevents knowledge leakage, since
2347 earlier KG snapshots do not contain information from later ones.

2348 **Are there any errors, sources of noise, or redundancies in the dataset?** We applied several
2349 quality-control measures, including removing duplicate or highly similar passages and filtering out
2350 passages with a low proportion of English words, among others described in Section 4.2. In addition,
2351 we manually annotated and verified a random subset of the dataset (see Section 4.2). Nevertheless,
2352 we do not consider EMERGE as entirely error-free, as it may contain factual inaccuracies resulting
2353 from erroneous edits in Wikipedia or Wikidata. Finally, the annotation agreement scores between the
2354 LLMs and human annotators, as well as between humans, are very strong (see Section Section 4.2)
2355 but not perfect, reflecting the complexity and intricacy of error detection in the dataset.

2356 **Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g.,
2357 websites, tweets, other datasets)?** Yes, the introduced EMERGE dataset is self-contained and
2358 consists of:

2359 1. Annotated instances containing passages with associated KG triples and TKGU operations.
2360 2. Wikidata KG snapshots to which the annotated TKGU updates are applied.

2361 **Does the dataset contain data that might be considered confidential (e.g., data that is pro-
2362 tected by legal privilege or by doctor–patient confidentiality, data that includes the content of
2363 individuals’ non-public communications)?** No, Wikidata and Wikipedia are public resources.

2364 **Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening,
2365 or might otherwise cause anxiety?** No, no such instances were observed in EMERGE.

2366 **Does the dataset identify any subpopulations (e.g., by age, gender)?** While Wikipedia and
2367 Wikidata contain entities from various subpopulations, when building EMERGE, we do not focus
2368 on identifying and annotating any one in particular.

2376 **Is it possible to identify individuals (that is, one or more natural persons), either directly or**
2377 **indirectly (that is, in combination with other data) from the dataset?** It is possible to identify
2378 individuals publicly described in Wikipedia pages or represented in Wikidata entities. However,
2379 we do not save other personal information, such as details of the editors involved in Wikipedia and
2380 Wikidata updates.

2381
2382 **Does the dataset contain data that might be considered sensitive in any way (e.g., data that**
2383 **reveals race or ethnic origins, sexual orientations, religious beliefs, political opinions or union**
2384 **memberships, or locations; financial or health data; biometric or genetic data; forms of gov-**
2385 **ernment identification, such as social security numbers; criminal history)?** Since Wikipedia
2386 and Wikidata are public resources intended to be factual, this concern can be disregarded for the
2387 majority of instances in EMERGE.

2388
2389 **L.3 COLLECTION PROCESS**

2390
2391 **How was the data associated with each instance acquired? Was the data directly observ-**
2392 **able (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indi-**
2393 **rectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for**
2394 **age or language)?** The EMERGE dataset was annotated using publicly available entity mentions
2395 in Wikipedia pages, as described in Section 4.1. These hyperlinked mentions are visible to any
2396 Wikipedia visitor as links to other pages. To annotate the TKGU operations, we relied on actual
2397 updates in Wikidata. Generative models (i.e., LLMs) were used only to verify whether the detected
2398 TKGU operations are reflected in the textual content of the passages (see Section 4.2).

2399
2400 **What mechanisms or procedures were used to collect the data (e.g., hardware apparatuses or**
2401 **sensors, manual human curation, software programs, software APIs)?** The EMERGE dataset
2402 was generated from the Wikipedia and Wikidata dumps of March 2025. A computing cluster with 64
2403 CPUs and 128 GB of RAM was used to process and generate the dataset. Additionally, a cluster with
2404 4 H100 GPUs was used to run Meta-Llama-3.1-405B for verifying that the TKGU operations
2405 are effectively represented in the textual passages (see Section 4.2).

2406
2407 **If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic,**
2408 **probabilistic with specific sampling probabilities)?** The test set used in our experiments was
2409 randomly sampled from the larger dataset, with a maximum of 1,000 instances per snapshot per
2410 KG delta. The sampling procedure, described in detail in Section 5, includes retention of
2411 a minimum of 400 instances per delta for operations that require actual updates to the KG (i.e., D-
2412 Triples, E-Triples, EE-Triples, and EE-KG-Triples). This ensures that the models are evaluated on
2413 a sufficiently large number of such instances. This is particularly important for D-Triples TKGU
2414 operations, which are very scarce in the original dataset; without this retention, a purely random
2415 subsample would contain only a few instances, potentially leading to high variability in the results.

2416
2417 **Who was involved in the data collection process (e.g., students, crowdworkers, contractors)**
2418 **and how were they compensated (e.g., how much were crowdworkers paid)?** The dataset was
2419 generated automatically from real-world updates to Wikidata and changes in Wikipedia articles.
2420 LLMs were used to assess each TKGU operation with respect to the knowledge contained in the
2421 textual passages. The only human involvement was the annotation of a subsample of the dataset to
2422 measure agreement with the LLM annotations. For this purpose, two researchers acted as annotators
2423 and were credited as co-authors of the paper.

2424
2425 **Over what timeframe was the data collected?** The data were collected from seven yearly snap-
2426 shots, spanning January 1, 2019, to January 1, 2025. For each snapshot, KG deltas were extracted
2427 for up to five weeks, ending on February 5 of the corresponding year.

2428
2429 **Were any ethical review processes conducted (e.g., by an institutional review board)?** No, the
public nature of the data, consisting of Wikipedia pages and Wikidata KG updates, meant that no
formal ethical review was required.

2430 **Did you collect the data from the individuals in question directly, or obtain it via third parties**
2431 **or other sources (e.g., websites)?** The data were obtained from publicly available Wikipedia and
2432 Wikidata repository dumps (<https://dumps.wikimedia.org/>).
2433

2434 **Were the individuals in question notified about the data collection?** No individuals were di-
2435 rectly involved in the data collection.
2436

2437 **Did the individuals in question consent to the collection and use of their data?** No individuals
2438 were directly involved in the data collection.
2439

2440 **If consent was obtained, were the consenting individuals provided with a mechanism to revoke**
2441 **their consent in the future or for certain uses?** No individuals were directly involved in the data
2442 collection.
2443

2444 **Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data**
2445 **protection impact analysis) been conducted?** No formal data protection impact analysis was
2446 conducted, as the dataset is derived entirely from publicly available Wikipedia pages and Wikidata
2447 KG updates and does not include private or sensitive information about individuals.
2448

2449 L.4 PREPROCESSING/CLEANING/LABELING

2450 **Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing,**
2451 **tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, process-**
2452 **ing of missing values)?** Yes. The original raw data from the Wikipedia and Wikidata dumps
2453 underwent several preprocessing steps:
2454

- 2455 1. Preprocessed Wikipedia wikitext, retaining only lists and textual paragraphs as dataset in-
2456 puts, while excluding tables, figures, and other multimodal elements.
2457
- 2458 2. Extracted only Wikipedia text containing explicitly annotated entity mentions by editors,
2459 which could be mapped to Wikidata updates within a given time window in the KG delta.
2460
- 2461 3. Constrained Wikipedia passages to lengths between 30 and 1,000 tokens.
2462
- 2463 4. Filtered out passages with fewer than 30% English words, using the Python `nltk` package.
2464
- 2465 5. Applied stability constraints by discarding changes in Wikidata and Wikipedia that were
2466 quickly rolled back (often indicating incorrect knowledge). Specifically, we retained Wiki-
2467 data KG updates persisting at least 7 days and Wikipedia edits not followed by another
2468 change within 30 minutes.
2469
- 2470 6. Ensured diversity by requiring passages aligned to similar updates in Wikipedia to differ
2471 in content, measured by edit distance (minimum 0.15 for texts under 2,500 characters and
2472 0.25 for texts 2,500 characters or longer).
2473
- 2474 7. Validated the alignment of TKGU operations to textual passages with LLMs, explicitly
2475 marking operations that could be grounded in the passage content (see Section 4.2 for
2476 further details).
2477

2478 **Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support**
2479 **unanticipated future uses)?** Yes. We preserved all input and output data from each preprocessing
2480 step, beginning with the raw Wikipedia and Wikidata dumps used to construct EMERGE.
2481

2482 **Is the software that was used to preprocess/clean/label the data available?** Yes, all the software
2483 that was used to preprocess/clean/label will be publicly released upon acceptance.
2484

2485 L.5 USES

2486 **Has the dataset been used for any tasks already?** Yes, in Section 5 we experiment with various
2487 current state-of-the-art information extraction models.
2488

2484 **Is there a repository that links to any or all papers or systems that use the dataset?** Yes,
2485 there is a repository (currently private due to anonymity policy), which will be made public upon
2486 acceptance.
2487

2488 **What (other) tasks could the dataset be used for?** Beyond the KG updating task presented in
2489 this paper, EMERGE could be directly applied to at least the following tasks:
2490

- 2491 1. Question answering over novel and emerging knowledge derived from the TKGU opera-
2492 tions introduced here.
- 2493 2. General knowledge graph completion, where certain changes may trigger additional up-
2494 dates that are not limited to entities mentioned in textual passages but instead depend on
2495 the evolving KG structure. To support this, we will release all KG changes, not only those
2496 aligned with textual passages, which form the core of EMERGE.
2497

2498 **Is there anything about the composition of the dataset or the way it was collected and prepro-
2499 cessed/cleaned/labeled that might impact future uses?** No.
2500

2501 **Are there tasks for which the dataset should not be used?** No.
2502

2503 L.6 DISTRIBUTION

2504 **Will the dataset be distributed to third parties outside of the entity (e.g., company, institu-
2505 tion, organization) on behalf of which the dataset was created?** Yes, the dataset will be made
2506 publicly available in Hugging Face.
2507

2508 **How will the dataset be distributed (e.g., tarball on website, API, GitHub)?** The
2509 EMERGE dataset will be distributed via Hugging Face (<https://huggingface.co/>), and
2510 the code for generating the dataset will be released on GitHub (<https://github.com/>).
2511

2512 **When will the dataset be distributed?** The EMERGE dataset will be released publicly upon
2513 acceptance of the paper.
2514

2515 **Will the dataset be distributed under a copyright or other intellectual property (IP) license,
2516 and/or under applicable terms of use (ToU)?** To support openness and collaboration in research,
2517 we release the datasets under the Creative Commons Attribution 4.0 International (CC BY 4.0)
2518 license. The full terms of this license can be found on the Creative Commons website: <https://creativecommons.org/licenses/by/4.0/>.
2519

2521 **Have any third parties imposed IP-based or other restrictions on the data associated with the
2522 instances?** No, the dataset is derived from publicly available Wikipedia and Wikidata knowledge
2523 repositories and is not subject to any third-party IP restrictions.
2524

2525 **Do any export controls or other regulatory restrictions apply to the dataset or to individual
2526 instances?** No, the dataset and its individual instances are based on publicly available Wikipedia
2527 and Wikidata content and are not subject to export controls or other regulatory restrictions.
2528

2529 L.7 MAINTENANCE

2531 **Who will be supporting/hosting/maintaining the dataset?** The dataset will be supported,
2532 hosted, and maintained by the authors of this paper.
2533

2534 **How can the owner/curator/manager of the dataset be contacted (e.g., email address)?** The
2535 dataset is curated and managed by the authors of this paper. Inquiries regarding the dataset, in-
2536 cluding access, usage, and reporting issues, can be directed to the corresponding authors via email.
2537 Additionally, users can submit questions or report issues through the GitHub repository hosting the
2538 dataset generation code.
2539

2538 **Is there an erratum?** No erratum has been issued for the EMERGE dataset.
2539

2540 **Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete in-**
2541 **stances)?** Yes, EMERGE will be regularly updated with emerging knowledge through yearly snap-
2542 **shots. Announcements regarding new versions will be communicated via the EMERGE GitHub**
2543 **repository. Additionally, as described in Section 4.4, users can generate customized versions of**
2544 **EMERGE by adjusting relevant hyperparameters, as well as personalized snapshots of different**
2545 **granularity (e.g., daily, weekly, monthly).**

2546 **If the dataset relates to people, are there applicable limits on the retention of the data asso-**
2547 **ciated with the instances (e.g., were the individuals in question told that their data would be**
2548 **retained for a fixed period of time and then deleted)?** The EMERGE dataset does not contain
2549 private or personally identifiable information about individuals. It is derived entirely from publicly
2550 available Wikipedia pages and Wikidata entities, and no retention limits for individual consent were
2551 applicable.

2552 **Will older versions of the dataset continue to be supported/hosted/maintained?** Yes, all pre-
2553 vious versions of EMERGE will continue to be supported, hosted, and maintained. Each version
2554 will be assigned a unique version number, and we will provide persistent links to access every ver-
2555 sion through Hugging Face storage server. This will ensure reproducibility of experiments and will
2556 enable users to reference or use specific dataset versions as needed.

2557 **If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for**
2558 **them to do so?** Yes. As described in Section 4.4, EMERGE users will have access to all necessary
2559 scripts to re-generate the dataset with customized settings. This includes adjusting hyperparameters
2560 such as the maximum passage length, generating the dataset for newer snapshots, and specifying the
2561 number and granularity of KG deltas.

2562 **M ACCESSIBILITY**

2563 The EMERGE will be released publicly via a Hugging Face repository. The accompanying code for
2564 extending it with emerging Wikipedia and Wikidata knowledge will be made available in a public
2565 GitHub repository. In addition, the test set used in our experiments is included as supplementary
2566 material with this submission.

2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591