ARTICLE IN PRESS

Information Processing and Management xxx (xxxx) xxx

Contents lists available at ScienceDirect

Information Processing and Management

journal homepage: www.elsevier.com/locate/ipm

Editorial

Large Language Models and Data Quality for Knowledge Graphs

Stefano Marchesin [a,*], Gianmaria Silvello [a], Omar Alonso [b,1]

[a] University of Padua, Padua, Italy
[b] Amazon, Palo Alto, CA, USA

ARTICLE INFO

Keywords:
Data quality
Knowledge graph
Large language model

ABSTRACT

Knowledge Graphs (KGs) have become essential for applications such as virtual assistants, web search, reasoning, and information access and management. Prominent examples include Wikidata, DBpedia, YAGO, and NELL, which large companies widely use for structuring and integrating data. Constructing KGs involves various AI-driven processes, including data integration, entity recognition, relation extraction, and active learning. However, automated methods often lead to sparsity and inaccuracies, making rigorous KG quality evaluation crucial for improving construction methodologies and ensuring reliable downstream applications. Despite its importance, large-scale KG quality assessment remains an underexplored research area.

The rise of Large Language Models (LLMs) introduces both opportunities and challenges for KG construction and evaluation. LLMs can enhance contextual understanding and reasoning in KG systems but also pose risks, such as introducing misinformation or "hallucinations" that could degrade KG integrity. Effectively integrating LLMs into KG workflows requires robust quality control mechanisms to manage errors and ensure trustworthiness.

This special issue explores the intersection of KGs and LLMs, emphasizing human–machine collaboration for KG construction and evaluation. We present contributions on LLM-assisted KG generation, large-scale KG quality assessment, and quality control mechanisms for mitigating LLM-induced errors. Topics covered include KG construction methodologies, LLM deployment in KG systems, scalable KG evaluation, human-in-the-loop approaches, domain-specific applications, and industrial KG maintenance. By advancing research in these areas, this issue fosters innovation at the convergence of KGs and LLMs.

1. Introduction

In recent years, Knowledge Graphs (KGs), which can be informally viewed as large, structured collections of relational facts, have emerged as central assets to support virtual assistants, search, and recommendations, especially on the Web. Notable examples are Wikidata, DBpedia, YAGO, and NELL. Moreover, KGs are increasingly used by companies to organize and reason over their data, with industry-scale KGs fusing data from various sources for downstream applications.

Building KGs lies at the intersection of data management and machine learning, encompassing techniques such as data integration, data cleaning, entity recognition and disambiguation, relation extraction, and active learning (Hogan et al., 2022; Weikum et al., 2021). However, these processes are largely automated and often imperfect, resulting in sparse or noisy knowledge graphs (KGs). As such, assessing the quality of KGs is crucial — not only to inform refinement of the construction pipeline, but also to maintain trustworthiness in applications that depend on accurate and comprehensive knowledge (Marchesin & Silvello, 2025).

* Corresponding author.
E-mail addresses: stefano.marchesin@unipd.it (S. Marchesin), gianmaria.silvello@unipd.it (G. Silvello), omralon@amazon.com (O. Alonso).

[1] Work does not relate to the author's position at Amazon.

https://doi.org/10.1016/j.ipm.2025.104281

Despite the centrality of this challenge, large-scale KG quality evaluation remains a relatively underexplored research area (Xue & Zou, 2023).

Concurrently, the emergence of Large Language Models (LLMs) offers both promising opportunities and significant challenges for the KG ecosystem (Dong, 2023; Pan et al., 2024). On the positive side, LLMs exhibit remarkable capabilities in semantic comprehension and contextual reasoning, which can boost KG development by filling in missing information, proposing relationships, or clarifying entities. Conversely, LLMs are susceptible to "hallucinations", outputs that are plausible but incorrect, and often fall short in maintaining factual accuracy, i.e., they struggle to "provide reliable information that is consistent with real-world knowledge" (Shami et al., 2025; Wang et al., 2024). This underscores a pressing need for reliability, credibility, and error management in the integration of LLMs into KG processes.

Motivated by these developments, this special issue promotes novel research on human–machine collaboration for KG construction and evaluation, fostering the intersection between KGs and LLMs. The rest of this editorial is as follows: Section 2 summarizes the contributions included in the special issue, and Section 3 provides concluding remarks.

2. Outline of the special issue

This special issue comprises 16 articles selected out of 47 submitted papers, the result of a rigorous peer-review process. Their contributions advance the state of the art and consider a wide range of emerging directions at the intersection of KGs and LLMs — including ontology engineering, commonsense and event knowledge acquisition, few-shot entity recognition, scalable KG validation, temporal and multimodal semantic reasoning, and human-aligned question generation. Collectively, these works highlight both methodological innovation and applications across various domains and knowledge-based systems.

LLM-Augmented Named Entity Recognition. Named Entity Recognition (NER) is a critical building block for KG construction, yet it remains a challenging task in low-resource and domain-specific environments where annotated data is scarce. Two articles in this special issue explore how LLMs can be strategically combined with prompt design, contrastive learning, and knowledge-aware fine-tuning to push the boundaries of NER performance in such settings.

Zhao, Gao, and Guo (2025), in their article "*A multi-granularity in-context learning method for few-shot Named Entity Recognition via Knowledgeable Parameters Fine-Tuning*", introduce MKIC, a novel framework designed to tackle the semantic and boundary recognition challenges in few-shot NER. Unlike standard prompt-based methods, MKIC leverages four auxiliary in-context learning tasks that capture fine-grained knowledge of entity labels, semantics, and boundaries. It also features a Knowledgeable Parameters Fine-tuning module, which enables targeted adaptation of the model without modifying the core LLM, preserving its generalization capabilities. Additionally, entity-level contrastive learning is applied to optimize inter- and intra-token distribution distances. MKIC demonstrates consistent improvements across eight datasets, outperforming prior few-shot NER approaches by up to 6.18% in F1 score, while remaining lightweight and efficient.

Wang et al. (2025) with their article entitled "*A novel large-language-model-driven framework for named entity recognition*" present LLMCC, a novel framework that reveals how different LLMs can complement each other when solving NER tasks. The authors introduce two new components – SemnRank, for filtering redundant demonstrations, and InforLaw-thought, for improving prompt informativeness – both of which contribute to more effective prompt engineering. LLMCC also incorporates an entity-aware contrastive learning strategy, further refining representation learning. Evaluated across five diverse domains, LLMCC significantly outperforms ten recent baselines, achieving over 5% gains in F1 score. Beyond performance, the work offers valuable insights into best practices for LLM selection, prompt design, demonstration selections, and training designs, marking an important advance in the integration of LLMs into NER and KG development pipelines.

Domain-Specific KG Construction. Four articles focus on building and validating KGs tailored to specialized domains – including medicine, science, cybersecurity, and legislation – where data quality is critical and domain-specific challenges demand tailored solutions. Across these contributions, LLMs play a pivotal role in knowledge extraction, enrichment, and validation.

Li et al. (2025), in their article "*Quality-Controllable automatic construction method of Chinese knowledge graph for medical decision-making applications*", address the unique challenges of building high-quality medical KGs in Chinese healthcare. Existing medical KGs often suffer from limited coverage, insufficient granularity, and inconsistent quality, reducing their utility in clinical applications. To overcome these limitations, the authors propose a set of construction principles informed by medical professionals, along with an automated pipeline that integrates chain-of-thought-based knowledge mining and axiom logic-based quality control. The resulting KG, WiMedKG, captures a wide spectrum of both commonsense and experiential medical knowledge, spanning 111 medical departments, 29 entity types, 128 relationship types, and 40 attribute types, while comprising over 350K entities, 3M relational triples, and 1M attribute triples. The KG has been assessed by medical professionals, achieving a validation score of 90.66%. Additionally, the authors show that enhancing an LLM with WiMedKG leads to a 1.51% average improvement on answering multiple-choice examination questions in different clinical occupations and career stages, reflecting the practical value of the developed KG.

In the domain of cybersecurity, the article "*PageLLM: Incremental approach for updating a Security Knowledge Graph by using Page ranking and Large language model*" by Mishra et al. (2025) tackles the challenge of maintaining high-quality KGs in a rapidly evolving threat landscape. Security KGs are inherently dynamic, and frequent retraining is computationally expensive. To address this, the authors introduce PageLLM, a method that combines PageRank-based importance scoring with LLM-driven active learning to update only the most impactful parts of the KG. This approach leverages LLMs to enrich contextual knowledge while ensuring scalability, allowing defenders to better anticipate and respond to emerging cyber threats.

The article "*An LLM-assisted ETL pipeline to build a high-quality knowledge graph of the Italian legislation*" by Colombo et al. (2025) explores the transformation of complex legal texts into structured, queryable KGs. Legislative systems are increasingly modeled using XML-based standards like Akoma Ntoso (AKN), but the resulting KGs often lack the depth needed for advanced legal analytics. To fill this gap, the authors present an Extract-Transform-Load (ETL) pipeline that integrates the AKN standard, implemented in the Italian system, with a property graph model, enabling advanced pattern detection. Fine-tuned LLMs are then used to enrich nodes and edges with high-value metadata. The result is a richly annotated legal KG that supports in-depth exploration of the Italian legislation.

Gajo and Barrón-Cedeño (2025), in "*Natural vs Programming Language in LLM Knowledge Graph Construction*", investigate the potential of lightweight, open-weight LLMs for constructing KGs from raw text, challenging the dominance of closed-source systems. Building on the CodeKGC prompting framework originally based on GPT-3.5, the authors evaluate LLMs pre-trained on either natural language or code, and fine-tuned using prompts formatted in both natural language and Python-like scripts. The findings reveal that prompt modality has minimal impact on performance and that chain-of-thought augmentation may actually degrade performance. Most notably, the considered lightweight LLMs outperform the much larger CodeKGC on both domain-specific (medicine, science) and general (news) datasets, underscoring the potential of small, open-weight LLMs to drive scalable and reproducible KG construction without reliance on proprietary APIs.

KG Validation and Quality Assessment. As KGs scale in size and scope, ensuring the accuracy and trustworthiness of their content becomes increasingly critical. Two articles in this special issue leverage LLMs – both independently and in human-in-the-loop systems – for validating KG triples at scale, moving beyond static rule-based checks and adopting context-aware verification strategies.

Adam and Kliegr (2025), in "*Traceable LLM-based validation of statements in knowledge graphs*", tackle a well-known shortcoming of LLMs: the inability to trace the source of their claims. To address this, the proposed method avoids relying on the LLM's internal knowledge, instead comparing KG statements against evidence retrieved through web searches or Wikipedia in a Retrieval-Augmented Generation (RAG) workflow. The method was tested on bioscience data from the BioRED dataset, achieving 88% precision but only 44% recall, underlining the need for human oversight. A separate evaluation using the SNLI dataset showed competitive results with models specialized in natural language inference. The approach was also applied to Wikidata, using SPARQL to extract candidate triples for validation. Overall, the method illustrates the potential of LLMs – when carefully grounded in external sources – for scalable and traceable KG verification.

Tsaneva et al. (2025), in their article "*Knowledge Graph Validation by Integrating LLMs and Human-in-the-Loop*", investigate various collaborative workflows that combine LLMs and human expertise. Tested within the construction pipeline of CS-KG, a massive Computer Science KG consisting of 350 million triples, the authors explore nine validation strategies ranging from fully automated to hybrid human–AI systems. While fully automated LLM validation boosts precision by 12%; however, it suffers from reduced recall, ultimately lowering the F1 score. In contrast, a hybrid human-in-the-loop approach delivers a 5% gain in F1, balancing quality and efficiency with minimal expert input. The study provides compelling evidence that LLMs can be integrated into KG validation pipelines not just as validators but as collaborative agents, helping to mitigate the bottlenecks of manual annotation at scale.

LLM-assisted Ontology Development. Ontologies serve as the backbone of KGs, providing the formal structure needed to ensure semantic coherence and enforce quality standards. Two articles in this special issue demonstrate how LLMs can facilitate the creation and refinement of ontologies, thereby reducing manual effort while enhancing scalability and consistency.

Val-Calvo et al. (2025), in their article "*OntoGenix: Leveraging Large Language Models for enhanced ontology engineering from datasets*", address the high cost and complexity of developing ontologies, particularly in enterprise contexts where datasets are vast, heterogeneous, and often unfamiliar to ontology engineers. The authors propose a structured, LLM-driven workflow – implemented in the OntoGenix tool – that supports key phases of ontology development: from data preprocessing and ontology planning to concept modeling and entity refinement. Evaluated across six commercial datasets, OntoGenix demonstrates the ability to produce coherent and human-like ontologies. Although human-created ontologies remain superior in handling the most complex modeling cases, OntoGenix offers a promising step towards semi-automated, scalable ontology development with LLMs as collaborative agents.

The article "*Large Language Models for Scholarly Ontology Generation: An Extensive Analysis in the Engineering Field*" by Aggarwal et al. (2025) investigates the use of LLMs for supporting the automatic generation of scholarly ontologies, focusing on semantic relation identification between research topics in the engineering domain. To this end, the authors introduce a novel benchmark derived from the IEEE Thesaurus, designed to evaluate model performance in identifying broader, narrower, and same-as relations. Seventeen LLMs are assessed, differing in scale, accessibility (open vs. proprietary), and architecture (full vs. quantized), across four zero-shot prompting strategies. Results show that full models achieve strong performance (F1 scores up to 0.967), while smaller quantized models, when paired with carefully designed prompts, can reach comparable accuracy with significantly lower computational cost. The study highlights the potential of LLMs – in particular of lightweight, open-weight variants – for scalable, high-quality ontology construction in scientific domains.

Temporal and Semantic Reasoning in KGs. Integrating temporal awareness and multimodal semantics into KGs represents a key challenge for evolving reasoning systems. Two articles in this special issue explore the use of LLMs and neuro-symbolic methods to enrich KGs with temporal structure and grounded meaning.

Xu et al. (2025), in "*Historical facts learning from Long–Short Terms with Language Model for Temporal Knowledge Graph Reasoning*", address the temporal dimension of KG reasoning, where the goal is to predict missing links by leveraging historical facts across time. Traditional GNN-based approaches and earlier LLM-based methods fall short in three areas: they fail to capture relational associations, lack nuanced treatment of temporal granularity, and often rely on external knowledge. The proposed Historical

Facts Learning (HFL) framework introduces timestamp tokenization and a multi-perspective sampling strategy to model the influence of past events better. HFL operates without external data, utilizing only the intrinsic content of the temporal KG or its textual representation. Tested across four Temporal KG Reasoning benchmarks, HFL demonstrates superior performance over both GNN-based and LLM-enhanced approaches. A variant adapted for LLMs further confirms the approach's flexibility and scalability.

De Giorgis et al. (2025), in their article "*Neurosymbolic graph enrichment for Grounded World Models*", propose a hybrid framework to fuse multimodal understanding with structured semantics. Starting from image inputs, their method utilizes LLMs to generate natural language descriptions, which are then transformed into Abstract Meaning Representation (AMR) graphs. These graphs are enriched with logical patterns and domain knowledge, including layered semantics derived from linguistic and factual KGs, and then fed back into the LLM for further inference. The enriched representation allows the system to reason over semantic implications, moral values, metaphors, and other context-dependent knowledge elements. This neuro-symbolic loop aims to bridge the gap between free-form LLM reasoning and formal symbolic representation, making it possible to construct grounded, interpretable world models that are better equipped for complex natural language understanding and high-stakes decision-making.

Commonsense and Event Knowledge. Commonsense reasoning and event understanding are foundational capabilities for intelligent systems. Three articles in this special issue present innovative approaches that leverage LLMs, synthetic data generation, and graph-based learning to enrich commonsense and event-centric knowledge in KGs. These works address persistent challenges, including data scarcity, relational depth, and representation quality.

The article "*Estimating the plausibility of commonsense statements by novelly fusing large language model and graph neural network*" by Yu, Lei, et al. (2025) proposes a hybrid approach to evaluate how plausible a given commonsense statement is – a central task in commonsense reasoning. The method combines the world knowledge encoded in LLMs with the structural expressiveness of Graph Neural Networks (GNNs). A novel maximum-paths mechanism is introduced to sample relevant subgraphs from the KG, which are embedded alongside LLM outputs into a multidimensional space, enabling richer reasoning. Experimental results across multiple benchmarks show the approach outperforms state-of-the-art methods, highlighting the value of LLM-GNN fusion in capturing nuanced commonsense patterns.

Yu, Tian, et al. (2025), in "*Amplifying commonsense knowledge via bi-directional relation integrated graph-based contrastive pre-training from large language models*", present BRIGHT, a bi-directional relation-integrated contrastive learning framework to improve commonsense KG acquisition. Unlike traditional methods that rely on unidirectional relations and costly human annotations, BRIGHT leverages LLMs to generate high-quality, diverse knowledge. By transforming relations into sentence-level bi-directional pairs and using contrastive pre-training with multiple negative samples, BRIGHT addresses the "reversal curse" and enhances semantic learning. It generates up to 397K novel commonsense triples and achieves top-1 accuracy rates of 90.51% on ATOMIC and 85.59% on ConceptNet KGs, respectively, thereby approaching human performance and significantly expanding existing commonsense resources.

The article "*Improving event representation learning via generating and utilizing synthetic data*" by Feng et al. (2025) tackles limitations in existing contrastive learning approaches for event modeling. Traditional methods rely heavily on dropout for augmentation and assume uniform similarity across all positive pairs, leading to coarse-grained alignment. The proposed LLM-CL framework introduces an LLM-driven, event KG-augmented synthetic data generation technique that produces semantically equivalent event pairs with varied text lengths and low lexical overlap. It also employs a self-adaptive contrastive loss that dynamically scales based on encoder-aware similarities. The approach achieves superior results in both intrinsic and downstream tasks, underscoring the potential of synthetic data and adaptive contrastive learning in enhancing event-centric KGs.

Human-Centric Question Generation from KGs. One article focuses on aligning KG-based question generation with human intent and quality through Reinforcement Learning from Human Feedback (RLHF).

Zhao, Tang, et al. (2025), in "*Towards human-like questioning: Knowledge base question generation with bias-corrected reinforcement learning from human feedback*", introduce BC-RLHF, the first framework to apply RLHF to the task of Knowledge Base Question Generation (KBQG). Existing KBQG models often produce fluent questions but lack diversity, authenticity, and a truly human questioning style. BC-RLHF addresses these gaps by first training a base question generation model and then developing two quality judgment models: one focused on human-likeness and the other on multi-granular quality aspects (e.g., diversity, informativeness, and relevance). These judges are integrated into a bias-aware RLHF pipeline designed to mitigate common issues such as sycophancy and confirmation bias during training. The result is a generation model capable of crafting more natural, engaging, and intent-aligned questions from structured KG input. Empirical results show significant improvements on multiple KBQG benchmarks, underscoring the impact of human-aligned optimization strategies for KG-to-text generation.

3. Conclusions

The contributions in this special issue offer a comprehensive and timely snapshot of the evolving relationship between KGs and LLMs. As the field moves beyond traditional graph-centric approaches, LLMs are reshaping how we design, construct, validate, and reason with structured knowledge. From enhancing ontology development workflows to scaling KG validation with human-in-the-loop LLM assistance, and from improving commonsense reasoning to generating more authentic, human-aligned questions, these articles pursue exciting directions towards the development of next-generation intelligent information access systems. At the same time, they also underscore important challenges: ensuring factual reliability, mitigating bias, maintaining explainability, and aligning automated outputs with domain-specific needs. Together, they reflect a research landscape that is not only technically dynamic but also increasingly interdisciplinary and human-centered. We hope this special issue will serve as both a reference and a catalyst for further exploration at this exciting frontier — where symbolic structure meets linguistic depth, and where KGs and LLMs converge to build more intelligent, reliable, and context-aware knowledge-based systems.

Acknowledgments

We thank all the authors for considering this special issue as a venue to publish their research at the intersection of LLMs and data quality for KGs. We are also grateful to the reviewers for their valuable and constructive feedback, which significantly contributed to the quality of the accepted articles. Finally, we extend our appreciation to the Editor-in-Chief and the editorial team for their guidance and support throughout the editorial process of this special issue.

References

Adam, D., & Kliegr, T. (2025). Traceable LLM-based validation of statements in knowledge graphs. *Information Processing & Management*, *62*(4), Article 104128. http://dx.doi.org/10.1016/j.ipm.2025.104128.

Aggarwal, T., Salatino, A., Osborne, F., & Motta, E. (2025). Large language models for scholarly ontology generation: An extensive analysis in the engineering field. *Information Processing & Management*.

Colombo, A., Bernasconi, A., & Ceri, S. (2025). An LLM-assisted ETL pipeline to build a high-quality knowledge graph of the Italian legislation. *Information Processing & Management*, *62*(4), Article 104082. http://dx.doi.org/10.1016/j.ipm.2025.104082.

De Giorgis, S., Gangemi, A., & Russo, A. (2025). Neurosymbolic graph enrichment for grounded world models. *Information Processing & Management*, *62*(4), Article 104127. http://dx.doi.org/10.1016/j.ipm.2025.104127.

Dong, X. L. (2023). Generations of knowledge graphs: The crazy ideas and the business impact. *Proceedings of the VLDB Endowment*, *16*(12), 4130–4137. http://dx.doi.org/10.14778/3611540.3611636.

Feng, Y., Li, L., Qin, X., & Zhang, B. (2025). Improving event representation learning via generating and utilizing synthetic data. *Information Processing & Management*, *62*(4), Article 104083. http://dx.doi.org/10.1016/j.ipm.2025.104083.

Gajo, P., & Barrón-Cedeño, A. (2025). Natural vs programming language in LLM knowledge graph construction. *Information Processing & Management*, *62*(5), Article 104195. http://dx.doi.org/10.1016/j.ipm.2025.104195.

Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutierrez, C., Kirrane, S., Labra Gayo, J. E., Navigli, R., Neumaier, S., Ngomo, A. C. N., P., A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J. F., Staab, S., & Zimmermann, A. (2022). Knowledge graphs. *ACM Computing Surveys*, *54*(4), 71:1–71:37. http://dx.doi.org/10.1145/3447772.

Li, X., Yuan, Y., Yang, Y., Guan, Y., Wang, H., Jiang, J., Shi, H., & Liu, X. (2025). Quality-Controllable automatic construction method of Chinese knowledge graph for medical decision-making applications. *Information Processing & Management*, *62*(4), Article 104148. http://dx.doi.org/10.1016/j.ipm.2025.104148.

Marchesin, S., & Silvello, G. (2025). Credible intervals for knowledge graph accuracy estimation. In *Proc. ACM manag. data 3, 3 (SIGMOD), article 142*. http://dx.doi.org/10.1145/3725279.

Mishra, C., Sarma, H., & M., S. (2025). PageLLM: Incremental approach for updating a security knowledge graph by using page ranking and large language model. *Information Processing & Management*, *62*(3), Article 104045. http://dx.doi.org/10.1016/j.ipm.2024.104045.

Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., & Wu, X. (2024). Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*, *36*(7), 3580–3599. http://dx.doi.org/10.1109/TKDE.2024.3352100.

Shami, F., Marchesin, S., & Silvello, G. (2025). Fact verification in knowledge graphs using LLMs. In *Proc. of the 48th international ACM SIGIR conference on research and development in information retrieval*. http://dx.doi.org/10.1145/3726302.3730142.

Tsaneva, S., Dessì, D., Osborne, F., & Sabou, M. (2025). Knowledge graph validation by integrating LLMs and human-in-the-loop. *Information Processing & Management*, *62*(5), Article 104145. http://dx.doi.org/10.1016/j.ipm.2025.104145.

Val-Calvo, M., Aranguren, M. E., Mulero-Hernández, J., Almagro-Hernández, G., Deshmukh, P., Bernabé-Díaz, J. A., Espinoza-Arias, P., Sánchez-Fernández, J. L., Mueller, J., & Fernández-Breis, J. T. (2025). OntoGenix: Leveraging large language models for enhanced ontology engineering from datasets. *Information Processing & Management*, *62*(3), Article 104042. http://dx.doi.org/10.1016/j.ipm.2024.104042.

Wang, Z., Chen, H., Xu, G., & Ren, M. (2025). A novel large-language-model-driven framework for named entity recognition. *Information Processing & Management*, *62*(3), Article 104054. http://dx.doi.org/10.1016/j.ipm.2024.104054.

Wang, Y., Wang, M., Manzoor, M. A., Liu, F., Georgiev, G. N., Das, R. J., & Nakov, P. (2024). Factuality of large language models: A survey. In *Proc. of the 2024 conference on empirical methods in natural language processing* (pp. 19519–19529). Association for Computational Linguistics, http://dx.doi.org/10.18653/v1/2024.emnlp-main.1088, URL https://aclanthology.org/2024.emnlp-main.1088/.

Weikum, G., Dong, X. L., Razniewski, S., & Suchanek, F. M. (2021). Machine knowledge: Creation and curation of comprehensive knowledge bases. *Foundation and Trends® in Databases*, *10*(2–4), 108–490. http://dx.doi.org/10.1561/1900000064.

Xu, W., Liu, B., Peng, M., Jiang, Z., Jia, X., Liu, K., Liu, L., & Peng, M. (2025). Historical facts learning from long-short terms with language model for temporal knowledge graph reasoning. *Information Processing & Management*, *62*(3), Article 104047. http://dx.doi.org/10.1016/j.ipm.2024.104047.

Xue, B., & Zou, L. (2023). Knowledge graph quality management: A comprehensive survey. *IEEE Transactions on Knowledge and Data Engineering*, *35*(5), 4969–4988. http://dx.doi.org/10.1109/TKDE.2022.3150080.

Yu, H. T., Lei, C., Ge, Y., Duan, Y., Liu, X., Lynden, S., Kim, K., Matono, A., & Jatowt, A. (2025). Estimating the plausibility of commonsense statements by novelly fusing large language model and graph neural network. *Information Processing & Management*, *62*(4), Article 104146. http://dx.doi.org/10.1016/j.ipm.2025.104146.

Yu, L., Tian, F., Kuang, P., & Zhou, F. (2025). Amplifying commonsense knowledge via bi-directional relation integrated graph-based contrastive pre-training from large language models. *Information Processing & Management*, *62*(3), Article 104068. http://dx.doi.org/10.1016/j.ipm.2025.104068.

Zhao, Q., Gao, T., & Guo, N. (2025). A multi-granularity in-context learning method for few-shot named entity recognition via knowledgeable parameters fine-tuning. *Information Processing & Management*, *62*(4), Article 104129. http://dx.doi.org/10.1016/j.ipm.2025.104129.

Zhao, R., Tang, J., Zeng, W., Guo, Y., & Zhao, X. (2025). Towards human-like questioning: Knowledge base question generation with bias-corrected reinforcement learning from human feedback. *Information Processing & Management*, *62*(3), Article 104044. http://dx.doi.org/10.1016/j.ipm.2024.104044.