



Contents lists available at ScienceDirect

Web Semantics: Science, Services and Agents on the World Wide Web

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

Logic Augmented Generation

Aldo Gangemi^{a,b,1}, Andrea Giovanni Nuzzolese^{b,*,1}^a University of Bologna, Bologna, Italy^b CNR - Institute of Cognitive Sciences and Technologies, Bologna, Italy

ARTICLE INFO

Keywords:

Knowledge graphs
Large language models
Logic augmented generation

ABSTRACT

Semantic Knowledge Graphs (SKG) face challenges with scalability, flexibility, contextual understanding, and handling unstructured or ambiguous information. However, they offer formal and structured knowledge enabling highly interpretable and reliable results by means of reasoning and querying. Large Language Models (LLMs) may overcome those limitations, making them suitable in open-ended tasks and unstructured environments. Nevertheless, LLMs are hardly interpretable and often unreliable. To take the best out of LLMs and SKGs, we envision Logic Augmented Generation (LAG) to combine the benefits of the two worlds. LAG uses LLMs as Reactive Continuous Knowledge Graphs that can generate potentially infinite relations and tacit knowledge on-demand. LAG uses SKGs to inject a discrete heuristic dimension with clear logical and factual boundaries. We exemplify LAG in two tasks of collective intelligence, i.e., medical diagnostics and climate projections. Understanding the properties and limitations of LAG, which are still mostly unknown, is of utmost importance for enabling a variety of tasks involving tacit knowledge in order to provide interpretable and effective results.

1. Introduction

Although there is no general consensus on the exact definition of Knowledge Graph (KG) in literature [1], we follow [2], and define KGs as linked data (typically expressed in the Web Ontology Language — OWL — and the Resource Description Framework — RDF triples) including both schema and assertional axioms. We refer to these KGs as Semantic Knowledge Graphs (SKGs), because of their rigorous model-theoretical semantics. SKGs provide structured, interpretable knowledge with logical querying and reasoning capabilities, making them valuable tools for modelling rich, linked datasets in any domain. However, their limitations in scalability, flexibility, and the integration of non-standard input restrict their application in open-ended tasks requiring complex extraction of perspectival and incomplete knowledge. *Collective intelligence* [3], as used in medical diagnostics and climate services is such a case. Additionally, SKGs: (i) are typically static and cannot natively represent dynamic, evolving knowledge (e.g. they struggle with temporal reasoning), making it difficult to handle changes over time, such as tracking historical trends or future predictions;

(ii) have limited ability to represent uncertainty or probabilistic relationships (Fuzzy OWL is an example [4]), which are often crucial in real-world scenarios and open-ended tasks; (iii) do not provide built-in mechanics to handle conflicting information from multiple sources²; (iv) are short of implicit or common-sense knowledge, which humans rely on extensively, thus reducing their effectiveness in tasks requiring contextual understanding or implicit knowledge extraction; (v) excel at representing explicit knowledge but are inadequate for extracting or encoding tacit knowledge, which is often crucial for decision-making and problem-solving in complex domains. To address these challenges, we introduce *Logic Augmented Generation* (LAG), a novel paradigm that integrates SKGs with Large Language Models (LLMs) to leverage the strengths of both approaches. LAG conceptualises LLMs as potential *Reactive Continuous Knowledge Graphs* (RCKGs), which can dynamically adapt to diverse inputs by extending and contextualising SKGs that work as base models. SKGs ensure logical consistency, enforce factual boundaries, and foster interoperability, while LLMs process unstructured data and provide contextual insights and tacit knowledge on demand. This dual approach enhances interpretability and reliability,

* Corresponding author.

E-mail addresses: aldo.gangemi@cnr.it (A. Gangemi), andreagiovanni.nuzzolese@cnr.it (A.G. Nuzzolese).URLs: <https://istc.cnr.it/people/aldo-gangemi> (A. Gangemi), <https://istc.cnr.it/people/andrea-nuzzolese> (A.G. Nuzzolese).¹ This author contributed equally to this work.² This can lead to inconsistencies and difficulty when reconciling diverse expert contributions or contradictory data.

<https://doi.org/10.1016/j.websem.2024.100859>

Received 20 November 2024; Received in revised form 20 December 2024; Accepted 23 December 2024

Available online 18 January 2025

1570-8268/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

mitigating the lack of truth-theoretic semantics in LLM’s output. We expect that LAG methods enable effective collaboration among experts and support the co-creation of actionable insights in complex, evolving contexts such as climate adaptation and collective medical decision-making.

2. Use case and problem framing

Semantic Knowledge Graphs (SKGs) provide structured, interpretable knowledge and support precise, logical querying. They have been used for modelling rich linked datasets in a variety of domains, such as health, cultural heritage, social and climate sciences, etc. In the context of the HACID³ project [5], we designed and modelled an SKG to support experts in open-ended diagnostic tasks with a *collective intelligence* [3] approach in two different domains, i.e., medical diagnostics and climate services. Collective intelligence is the shared, tacit, or group intelligence that emerges from the collaboration, collective efforts, and interactions of a group of individuals or entities. This concept is rooted in the idea that a group whose members work together effectively can achieve insights, solve problems, or make decisions that are beyond the ability of a single individual.

We use the following simple medical case, referred to as CASE, as an example to clarify the problem.

“A 38-year-old male presents with fever and a dry cough that has persisted for the last 4 days. He has no significant medical history, although he recently returned from a business trip.”

In the medical case above, the business trip is the potential cause of the patient’s medical condition. However, the causality relation is inferred by recognising the potential tacit connection between the business trip and the illness (such as travel-related infections) without it being directly stated or formally encoded in the information provided. It also includes the ability of physicians to identify and weigh subtle clues or patterns that may not be immediately obvious from the presented data.

Consequently, the objective is to use an SKG to smoothen the collaboration and communication among experts by harmonising the diversity of thought in a context in which no single individual has all the answers, but the group as a whole possesses the knowledge needed to solve a problem. In medical diagnostics, the ultimate goal is to identify the correct diagnosis for a patient, thereby reducing errors, preserving life, and minimising the costs associated with incorrect treatments.

Climate services urgently require systems that leverage collective intelligence to navigate and integrate vast, diverse datasets. By fostering collaboration among scientists, stakeholders, and AI-driven tools, challenges in the interpretability of results and the handling of unstructured data can be addressed more effectively.

To address those objectives, our SKG integrates different existing knowledge bases and provides a coherent conceptualisation formalised in a rich modular ontology network designed by reusing Ontology Design Patterns [6] (ODPs) and leveraging DOLCE-Zero [7] as a top-level ontology. For medical diagnostics, we have integrated data gathered from SNOMED-CT, ICD-10, and Wikidata. For climate services, we have integrated data from the Coupled Model Intercomparison Project⁴ (CMIP) [8].

Nevertheless, the effectiveness of an SKG solely based on symbolic knowledge representation and reasoning remains limited when we need to support open-ended tasks with expansive problem spaces, such as collective intelligence-based medical diagnostics, or climate services. This limitation arises from the challenges of harmonising and

integrating standardised knowledge in the reference SKG with non-standard, natural language inputs from multiple experts. Firstly, the different inputs provided by experts might express different viewpoints in different languages and factual granularities. Secondly, tacit and domain-specific background knowledge might be required in order to make sense of those inputs, and making a superior collective knowledge emerge. Moreover, SKGs struggle with scalability, flexibility, evolvability, and the integration of complex, context-rich information. For example, climate data include real-time environmental observations, policy impacts, and social adaptation strategies.

Large Language Models (LLMs), in contrast, can be employed to manage unstructured data and to generate novel insights on demand, making them well-suited for open-ended tasks. Yet, their lack of interpretability and reliability hinders their application in fields where precise and accountable outputs are critical.

3. LLMs as reactive continuous knowledge graphs with logic boundaries

To address this gap, we use LLMs as potential *Reactive Continuous Knowledge Graphs* (RCKGs). RCKGs are SKGs extracted from multimodal signal (e.g. text, speech, pictures, sensory data, etc.), using (typically or mostly) continuous vector spaces, such as generative pre-trained transformers. RCKG extraction involves a three-step process, i.e. (i) mapping multimodal signal to natural language, (ii) converting natural language into an SKG, and (iii) extending the SKG with tacit knowledge based on multiple heuristics. Each transformation in this pipeline plays a crucial role in enabling the RCKG’s ability to generate dynamic, context-sensitive knowledge representations. The first transformation, mapping signal to natural language, is characterised as *supramodal* [9,10] because natural language inherently integrates and synthesises information from multiple sensory modalities (e.g., visual, auditory, tactile inputs) by representing objects and concepts that are invariant across different sensory modalities. For instance, a textual description of an image consolidates some visual details, or a written summary of a sound encapsulates certain auditory characteristics. Natural language serves as a unified, modality-transcending medium that bridges sensory-specific data and symbolic representation. This supramodal nature makes natural language an ideal intermediary for integrating multimodal signals into a coherent, human-readable format that captures both explicit information and implicit nuances, which are more or less easily reconstructed by human interpreters. The second transformation, converting natural language to an SKG, is defined as *amodal* [10] because knowledge graphs abstract knowledge from its sensory or linguistic origins. In this step, semantic relationships and entities are encoded in a structured, modality-neutral format that is fully machine-interpretable. This abstraction ensures that the resulting representation is independent of the modalities through which the original information was perceived or described. The amodal KG thus becomes a foundation for logical reasoning, enabling precise querying, inference, and integration with external knowledge sources, such as SKGs. The third transformation, extending the SKG with implicit/tacit knowledge, aims at recovering the tacit dimension of natural language as used in context, by leveraging reactive natural language models such as generative LLMs as providers. This extension results in Reactive Continuous Knowledge Graphs (RCKGs). The third transformation again exploits the supramodal nature of language to generate amodal knowledge.

The relevance of this supramodal-to-amodal pipeline in RCKG lies in its ability to leverage the strengths of both transformations. By systematically aligning supramodal and amodal transformations, a RCKG provides a scalable and versatile framework for synthesising, abstracting, and reasoning with multimodal inputs, thereby advancing the potential of neuro-symbolic systems in dynamic and open-ended tasks.

An SKG can be generated from text in multiple ways. State-of-the-art SKG extraction from text assumes *situations* as occurrences of conceptual *frames*, following the knowledge extraction paradigm introduced

³ The HACID acronym stands for Hybrid Human Artificial Collective Intelligence in Open-Ended Decision Making project.

⁴ <https://www.wcrp-climate.org/wgcm-cmip>.

by the FRED machine reader [11], which extracts OWL SKGs based on Framester Semantics [12], and is also used [13] to formalise Abstract Meaning Representation [14] (AMR) graphs. Once extracted, an SKG is passed (within an augmented prompt) to an LLM as a source of extended meaning, to contextualise and complete an SKG with implicit and local knowledge. In this sense, a LLM is used as a reactive AI, which generates a piece of extended knowledge K_{ext} when stimulated by an input piece of knowledge K_{inp} .

A first implementation of LAG as SKG-augmented RCKG generation is demoed as a multimodal knowledge extractor⁵ [15], used to generate the example graph in Fig. 2.

While the notions of reactivity and continuity in Knowledge Graphs (KGs) have established meanings in the literature, Reactive Continuous Knowledge Graphs (RCKGs) add complementary interpretations.

The reactivity of a Knowledge Graph (KG) has typically been associated with notions of dynamism and temporal awareness. A notable example is the work in [16], which explores the extension of KGs to accommodate temporal and dynamic aspects. Its authors focus on explicitly encoding time-aware data through timestamps or mappings across time points, thereby preserving historical and future contexts. More broadly, Reactive Continuous Knowledge Graphs (RCKGs) integrate tacit knowledge and dynamically adapt it using LLMs. RCKGs capture reasoning and contextual nuances implicit in data by generating new knowledge adaptively, blending probabilistic and logical reasoning in a manner that extends beyond static temporal annotations.

In addition to our definition of reactive and (mathematically) continuous KGs, Continuous Knowledge Graphs have been introduced recently, e.g. in [17] as evolutionary dynamic graphs, featuring incremental knowledge updates, schema adaptation, conflict resolution, source tracking and verification, and temporal awareness. The authors propose to reduce extraction to the normalisation and manipulation of elementary propositions, but recognise that this is a very challenging task. As said above, logic augmented generation is better suited to represent *situations* as occurrences of conceptual *frames* in a frame semantics-compatible logic like OWL2.

The notion of continuity proposed by [17] focuses on the process of updating dynamic and adaptable knowledge representation systems that can evolve and integrate new information systematically over time. The authors in [17] discuss the well-known task of continuously evolving graphs over time using available information. In addition, RCKGs are continuous firstly because LLMs are reactive continuous models: they are mostly trained on a finite and discrete corpus of data, which forms the basis of their learned knowledge; however, due to their ability to generalise from training data, they can generate a potentially infinite number of outputs by combining concepts, patterns and relationships in novel ways. This generative capability is what allows LLMs to learn and generate responses based on context-sensitive input data (such as text) and to continuously adapt to new contexts or queries during in-context learning. Hence, by means of prompting, their internal representation of knowledge adapts dynamically as more data is processed, allowing them to dynamically adjust and generate relevant responses based on the latest input. Accordingly, if properly prompted to generate a KG, a LLM can generate potentially infinite triples by means of in-context learning. This means, for example, that LLMs do not require retraining, adjustments or evolution to their underlying data to generate new knowledge. More specifically, this characteristic refers specifically to the ability of LLMs to contextually interpret and generate knowledge based on input prompts, rather than performing fundamental updates to their internal weights or training data. This adaptation is achieved through in-context learning, where the model dynamically adjusts its responses based on the structure and content of the input without altering its pre-trained parameters. LLMs with fixed weights rely on their generative capabilities, drawing from a latent

knowledge space encoded during training. While this space is finite and static, the models are capable of producing effectively infinite outputs by combining and contextualising learned patterns in novel ways. This makes them highly flexible for generating new knowledge dynamically, in response to specific prompts, without requiring retraining or modifications to the underlying knowledge corpus. In contrast, approaches like fine-tuning or continual learning involve fundamental updates to the model's weights, enabling the incorporation of new knowledge at the cost of additional computational overhead and potential domain-specific constraints. By focusing on prompting and in-context learning, LAG leverages on fixed-weight LLMs to serve as RCKGs, adapting dynamically to new inputs without requiring structural adjustments to the model.

4. LAG: SKG+RCKG

RCKGs are an example of a neuro-symbolic approach that generates a much richer and deeper knowledge graph than the basic meaning of the original text. RCKGs extract tacit knowledge by capturing insights, reasoning, or behaviours that resemble the informal, experiential, and often unspoken knowledge humans acquire through experience and practice [18]. While tacit knowledge is difficult to articulate or codify, it plays a crucial role in problem-solving and decision-making, making it particularly valuable for collective intelligence tasks that require collaboration, intuition, creativity, and experience. In medical diagnostics, tacit knowledge could involve a clinician's ability to identify subtle patterns that suggest a rare disease based on prior experience or an intuitive sense of urgency derived from non-verbal cues of a patient. For example, while symptoms such as fever and cough can be codified in structured medical records, a doctor's recognition that recent travel increases the likelihood of a tropical infection involves tacit knowledge. Similarly, in climate services, tacit knowledge includes understanding the implicit risks associated with certain weather patterns. For example, a meteorologist might intuitively anticipate that an unusual sea surface temperature anomaly could lead to intensified tropical storms, even before formal models confirm it. Similarly, local knowledge about how specific communities respond to extreme weather events is often unspoken, but vital for effective planning and adaptation strategies. Verifying and interpreting tacit knowledge is challenging, particularly in terms of reliability and ethical considerations, as it is not derived from human personal experience, validated sensors, or deductive inferences, but from patterns in data.

In the case of LLMs, tacit knowledge is implicitly encoded in the vast training data they process. LLMs do not *understand* tacit knowledge as humans do, but capture and generalise patterns, correlations, and relationships present in the data [19]. These models are trained on diverse multimodal corpora, which often include fragments of tacit knowledge embedded in human language, such as metaphors, cultural nuances, or context-specific inferences. This training phase does not explicitly codify tacit knowledge but instead creates a latent space where such patterns are represented implicitly as probabilistic associations across continuous vector spaces. Hence, we assume that the generative capabilities of LLMs allow them to be used as RCKG to make tacit knowledge more explicit when prompted effectively. Hence, by leveraging the ability of LLMs to synthesise patterns from training data, RCKG can generate representations (i.e. triples) that resemble the tacit insights humans infer from their experiences. For example, when presented with a prompt that implicitly contains causal relationships or contextual nuances, the model can produce explicit knowledge triples or statements that align with those relationships. This process can be understood as a "*reactive codification*", where tacit knowledge is externalised as explicit, structured outputs in response to specific inputs. In a nutshell, LAG exploits the duality between implicit tacit knowledge (i.e. the one generated during training) and explicit tacit knowledge (i.e. the one externalised in the form of structured knowledge graphs through generation and prompting). Accordingly, LAG use RCKGs to

⁵ <https://arco.istc.cnr.it/itaf/>

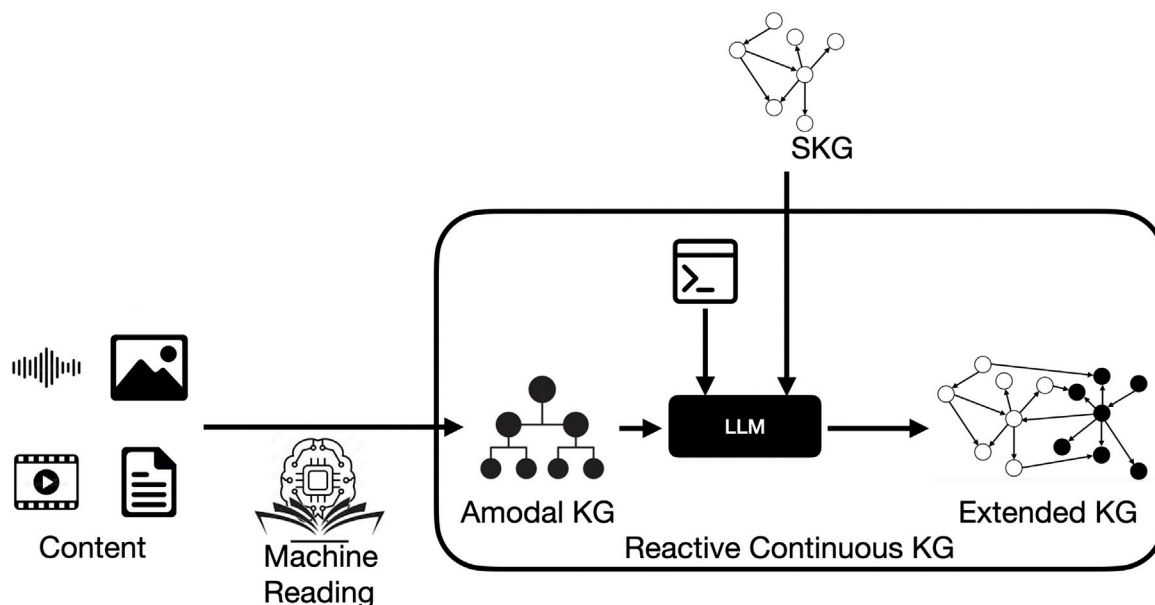


Fig. 1. The architecture of LAG.

surface tacit knowledge that would otherwise remain latent in the training data. The process hinges on prompt engineering and in-context learning, which guide the model to retrieve, synthesise, and represent tacit patterns relevant to a specific query or context. Nevertheless, a semantics for RCKG tacit knowledge requires addressing the nature of reactive inferences in a continuous space with discrete input and output that can be context-sensitive.

This puts RCKG apart from most logics. For example, RCKG semantics is *plausibility-preserving* rather than *truth-preserving*. Its operators for entailment, necessity, possibility, etc. would be interpreted differently from traditional truth-theoretic ones, and the interaction between classical truth-preserving and plausibility-preserving axioms needs to be negotiated. E.g., plausibility does not guarantee truth, while truth guarantees plausibility, leading to potential invalidity of mixed transitivity chains. Also, tacitness is typically non-monotonic and context-dependent. Furthermore, multiple tacit operators emerge which can vary in their reasoning behaviour: social implicatures can have different logical properties (group relativity, role sensitivity, status dependency, intensity, valence, gradability, source dependency, etc.) from affective evocations, evaluative judgements, causal consequences, metaphorical blending, etc. However, being able to represent such a variety of tacit knowledge by using a unique continuous model, and forcing it to adhere to the valid logical axioms of an amodal KG is a big step forward to use logic in practical, large scale, and flexible contexts that depend on sophisticated implicit knowledge.

Constraining RCKGs to generate knowledge within precise boundaries poses significant challenges. However, we deem the effort with formalising tacit knowledge valuable. This involves limiting the continuous geometric space to a discrete dimension, which allows, for instance, to get only the relevant knowledge with respect to a given objective (e.g., generating climate projections that harmonise experts' opinions and target expected environmental scenarios). To address these issues we started formalising and implementing a novel architecture called Logic Augmented Generation (LAG). LAG integrates RCKGs with Semantic KGs (SKGs) that serve as a discrete, heuristic layer for in-context learning within RCKG's continuous knowledge space, enforcing hybrid logical consistency, factual boundaries, and fostering interoperability with existing SKGs in the Semantic Web. Fig. 1 exemplifies the architecture of a LAG that is composed of: (i) an amodal SKG automatically gathered from text using FRED, representing user queries; (ii) an LLM that is prompted with the amodal SKG, additional

prompt heuristics that can be used to refine the context incrementally, and the SKG scoping the desired semantics and providing existing factual knowledge; (iii) an extended SKG produced by LAG, which contains newly generated triples that comply with the SKG, and addresses the desired user queries. An SKG is enriched as a RCKG via in-context learning that requires adequate prompt engineering strategies. The study of effective prompt engineering strategies to inject the logical and factual knowledge of an SKG into a RCKG is another challenge that LAG must address. Evolutions of Chain-of-Thought, such as Metacognitive Prompting [20], can be beneficial to LAG, as shown by preliminary results presented in [21]. LAG differs from existing paradigms available in literature such as Retrieval-Augmented Generation (RAG) [22] and fine-tuning, addressing critical gaps to enable dynamic and context-sensitive knowledge generation. In fact, RAG enhances LLMs by retrieving relevant information from a fixed corpus to ground responses in explicit factual knowledge. However, it is limited to static corpora and lacks mechanisms to generate or adapt knowledge dynamically. In contrast, LAG leverages RCKGs, which enable dynamic knowledge synthesis by reacting to input prompts, and integrates them with SKGs to ensure logical consistency and factual alignment. Unlike fine-tuning, which requires modifying model weights and retraining to incorporate new knowledge, LAG operates through in-context learning, injecting SKGs as structured prompts into RCKGs. This approach preserves the adaptability and domain neutrality of LLMs while enabling real-time, context-aware knowledge generation. Furthermore, while fine-tuning risks inconsistencies across domains and RAG relies on static retrieval, LAG supports reasoning over evolving knowledge spaces and accommodates tacit knowledge, probabilistic inferences, and hybrid logical reasoning. Additionally, by using a neuro-symbolic approach, LAG facilitates the integration of structured knowledge with generative reasoning, allowing it to address challenges such as handling dynamic and conflicting information. This positions LAG as a complementary and extended paradigm, particularly suited for applications requiring both knowledge generation and reconciliation, such as collective intelligence and collaborative decision-making.

An example that explains how LAG supports collective intelligence for medical diagnostics is based on the simple medical case we presented in Section 2, i.e. CASE. A possible extended KG generated by LAG for CASE is presented in Fig. 2. In this figure, the amodal KG generated through machine reading is represented by the black nodes and arrows, while the red ones are added to the extended

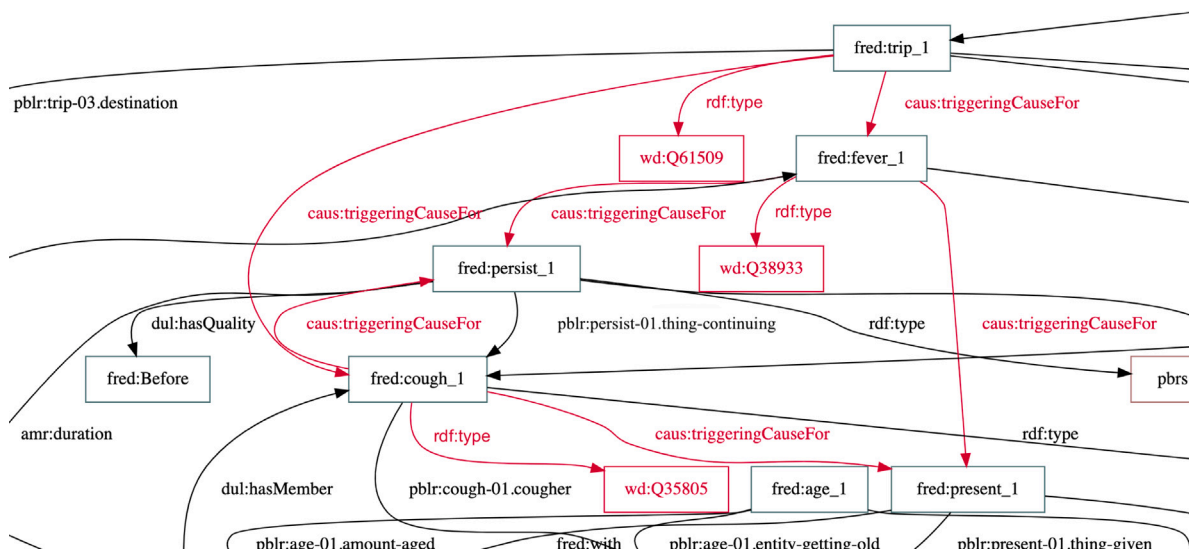


Fig. 2. Example of possible extended KG generated by LAG.

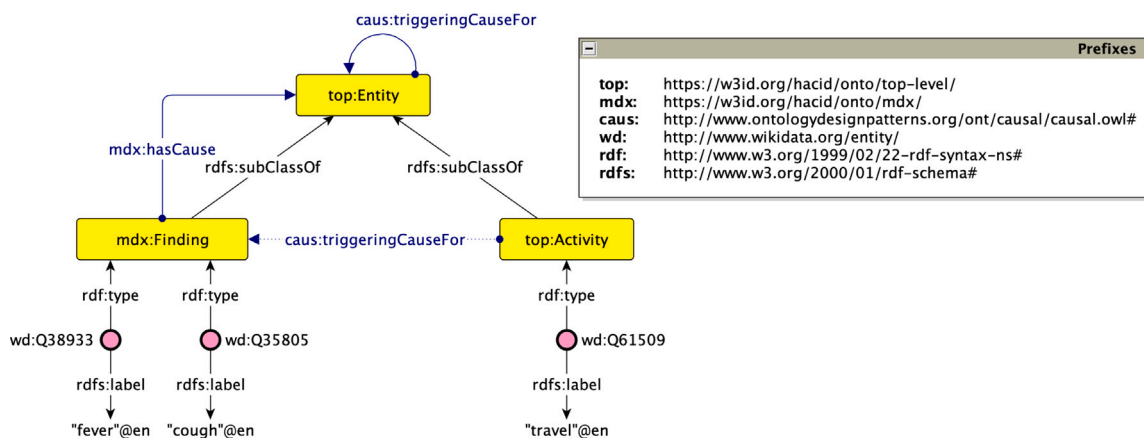


Fig. 3. A part of the SKG used as input of LAG to demonstrate our example (i.e. CASE).

KG by LAG. A part of the SKG used as input of LAG is depicted as a Grafoo⁶ [23] diagram in Fig. 3 that shows some core concepts with their semantics. Those concepts are: (i) top:Entity, which is the top level class in our ontology network; (ii) mdx:Finding, which is the class for representing any possible finding associated with medical relevance; and (iii) top:Activity, which represents any action or task planned or executed by an agent. The ontology network used for the medical case is described in a project deliverable [24] and can be retrieved on GitHub⁷. Additionally, Wikidata entities are used to populate the extensional layer of the SKG. Specifically, the entities wd:Q61509, wd:Q38933, and wd:Q35805 are sourced from Wikidata and represent the concepts travel, fever, and cough, respectively. The object properties caus:triggeringCauseFor and mdx:hasCause represent general and specific causal relationships, respectively. Indeed, the extended KG captures causal consequences generated as tacit knowledge by the RCKG component of LAG, i.e., the predicates caus:triggeringCauseFor, from the presented medical case. Then, the reconciliation with Wikidata performed by LAG involves the following steps: (i) entity matching and selection in which relevant entities from Wikidata are identified through entity

linking techniques. For instance, terms like “fever” or “cough” in the medical use case are mapped to their corresponding Wikidata entries (e.g., wd:Q38933 for fever), thus ensuring semantic alignment between the input data and the SKG; (ii) contextual-relevant knowledge gathering in which a subset of Wikidata is selected based on the contextual relevance of entities to the task or domain. For example, in the medical diagnostics use case, entities related to diseases, symptoms, and potential causes (such as travel history) are prioritised. This subset can be extracted, for example, using SPARQL queries that traverse relationships and filter entities based on predefined criteria (e.g., medical conditions with known causes or symptoms); (iii) knowledge harmonisation that allows the extracted subset of Wikidata to be incorporated into the extended KG. In the example in Fig. 2 the extended KG makes formally emerge that a trip is the potential triggering cause for both fever and cough, a cause that was only tacitly assumed in the case description. Hence, experts might rely on the extended KG in order to provide their opinions that can be aggregated to the extended KG following the same process. This approach has the potential to foster collaboration and improve decision-making in domains such as medical diagnostics and climate services, where the integration of multiple experts’ views, encompassing both tacit and formal knowledge, is crucial. It is worth noting that implementing LAG requires facing several challenges. First, SKG might contain broad, general-purpose knowledge, which requires filtering irrelevant or noisy data. With respect to this point, frame-based approaches have been demonstrated

⁶ <https://essepuntato.it/graffoo/>

⁷ <https://github.com/hacid-project/knowledge-graph/blob/main/ontologies/medical-dx/mdx.owl>

to be effective to get meaningful knowledge boundaries around broad data. An example is provided in [25], which uses knowledge patterns, a specific type of frame, to filter relevant contextual knowledge from DBpedia based on user queries. Another challenge to address is the reconciliation between the amodal KG and the SKG (e.g., Wikidata) to preserve the structure and specific schema of the latter by consistently mapping relationships and types. Finally, efficiently navigating and querying a large SKG (e.g., Wikidata) requires optimised query strategies and pre-computed indexes.

5. Conclusions and future work

In this work we introduce Logic Augmented Generation (LAG), a novel solution that integrates Semantic Knowledge Graphs (SKGs) with Reactive Continuous Knowledge Graphs (RCKGs) to enforce logical consistency and blend symbolic and continuous reasoning. RCKGs are products of a neuro-symbolic method that uses Large Language Models (LLMs) to extract and represent knowledge, incorporating tacit insights and contextual nuances. This hybrid model enables context-aware knowledge generation and addresses challenges like formalising tacit knowledge, enhancing prompt engineering, and aligning plausibility- and truth-preserving logic. The potentiality of LAG is illustrated by using collective intelligence as a case study. Future work will study SKG-RCKG integration, advance prompt strategies, and expand operators for representing tacit semantics.

CRedit authorship contribution statement

Aldo Gangemi: Writing – review & editing, Writing – original draft, Conceptualization. **Andrea Giovanni Nuzzolese:** Writing – review & editing, Writing – original draft, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the European Union's Horizon Europe research and innovation programme within the context of the project HACID (Hybrid Human Artificial Collective Intelligence in Open-Ended Domains, grant agreement No 101070588), and FAIR (Future Artificial Intelligence Research PNRR project, code PE00000013 CUP 53C22003630006)

Data availability

No data was used for the research described in the article.

References

- [1] S. Ji, S. Pan, E. Cambria, P. Marttinen, S.Y. Philip, A survey on knowledge graphs: Representation, acquisition, and applications, *IEEE Trans. Neural Netw. Learn. Syst.* 33 (2) (2021) 494–514.
- [2] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G.D. Melo, C. Gutierrez, S. Kirrane, J.E.L. Gayo, R. Navigli, S. Neumaier, et al., Knowledge graphs, *ACM Comput. Surv. (Csur)* 54 (4) (2021) 1–37.
- [3] J.M. Leimeister, Collective intelligence, *Bus. Inf. Syst. Eng.* 2 (2010) 245–248.
- [4] G. Stoilos, G.B. Stamou, V. Tzouvaras, J.Z. Pan, I. Horrocks, Fuzzy OWL: Uncertainty and the semantic web, in: *OWLED*, Vol. 5, 2005, pp. 11–12.
- [5] R.H. Kurvers, A.G. Nuzzolese, A. Russo, G. Barabucci, S.M. Herzog, V. Trianni, Automating hybrid collective intelligence in open-ended medical diagnostics, *Proc. Natl. Acad. Sci.* 120 (34) (2023) e2221473120.
- [6] A. Gangemi, V. Presutti, Ontology design patterns, in: *Handbook on Ontologies*, Springer, 2009, pp. 221–243.
- [7] H. Paulheim, A. Gangemi, Serving dbpedia with DOLCE—more than just adding a cherry on top, in: *The Semantic Web-ISWC 2015: 14th International Semantic Web Conference*, Bethlehem, PA, USA, October 11–15, 2015, Proceedings, Part I 14, Springer, 2015, pp. 180–196.
- [8] V. Eyring, S. Bony, G.A. Meehl, C.A. Senior, B. Stevens, R.J. Stouffer, K.E. Taylor, Overview of the coupled model intercomparison project phase 6 (CMIP6) experimental design and organization, *Geosci. Model Dev.* 9 (5) (2016) 1937–1958.
- [9] S.L. Fairhall, A. Caramazza, Brain regions that represent amodal conceptual knowledge, *J. Neurosci.* 33 (25) (2013) 10552–10558.
- [10] J.R. Binder, In defense of abstract conceptual representations, *Psychon. Bull. Rev.* 23 (4) (2016) 1096–1108.
- [11] A. Gangemi, V. Presutti, D. Reforgiato Recupero, A.G. Nuzzolese, F. Draicchio, M. Mongioli, Semantic web machine reading with FRED, *Semant. Web* 8 (6) (2017) 873–893.
- [12] A. Gangemi, M. Alam, L. Asprino, V. Presutti, D.R. Recupero, Framester: A wide coverage linguistic linked data hub, in: E. Blomqvist, P. Ciancarini, F. Poggi, F. Vitali (Eds.), *EKAW*, in: *Lecture Notes in Computer Science*, vol. 10024, 2016, pp. 239–254.
- [13] A. Gangemi, A. Graciotti, A. Meloni, A.G. Nuzzolese, V. Presutti, D. Reforgiato Recupero, A. Russo, R. Tripodi, Text2AMR2FRED, a tool for transforming text into RDF/OWL knowledge graphs via abstract meaning representation, in: *Proceedings of the ISWC 2023 Posters, Demos and Industry Tracks: from Novel Ideas to Industrial Practice Co-Located with 22nd International Semantic Web Conference*, ISWC 2023, 2023.
- [14] L. Banarescu, C. Bonial, S. Cai, M. Georgescu, K. Griffith, U. Hermjakob, K. Knight, P. Koehn, M. Palmer, N. Schneider, Abstract meaning representation for sembanking, in: *Proc. of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, ACL, Sofia, Bulgaria, 2013, pp. 178–186, <https://aclanthology.org/W13-2322>.
- [15] S. De Giorgis, A. Gangemi, A. Russo, Neurosymbolic graph enrichment for grounded world models, 2024, arXiv:2411.12671.
- [16] F. Krause, T. Weller, H. Paulheim, On a generalized framework for time-aware knowledge graphs, in: *Towards a Knowledge-Aware AI*, IOS Press, 2022, pp. 69–74.
- [17] Y. Lairgi, L. Moncla, R. Cazabet, K. Benabdeslem, P. Cléau, iText2KG: Incremental Knowledge Graphs Construction Using Large Language Models, 2024, arXiv:2409.03284.
- [18] M. Polanyi, The tacit dimension, in: *Knowledge in Organisations*, Routledge, 2009, pp. 135–146.
- [19] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, et al., A survey on evaluation of large language models, *ACM Trans. Intell. Syst. Technol.* 15 (3) (2024) 1–45.
- [20] Y. Wang, Y. Zhao, Metacognitive prompting improves understanding in large language models, 2023, arXiv preprint arXiv:2308.05342.
- [21] A.S. Lippolis, M. Ceriani, S. Zuppiroli, A.G. Nuzzolese, Ontogonia: Ontology generation with metacognitive prompting in large language models, in: *Proc. of the Extended Semantic Web Conference - Satellite Events*, Springer, Crete, Greece, 2024.
- [22] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive nlp tasks, *Adv. Neural Inf. Process. Syst.* 33 (2020) 9459–9474.
- [23] R. Falco, A. Gangemi, S. Peroni, D.M. Shotton, F. Vitali, Modelling OWL ontologies with Graffoo, in: V. Presutti, E. Blomqvist, R. Troncy, H. Sack, I. Papadakis, A. Tordai (Eds.), *ESWC (Satellite Events)*, in: *Lecture Notes in Computer Science*, vol. 8798, Springer, 2014, pp. 320–325.
- [24] A.G. Nuzzolese, A.S. Lippolis, W. Xu, M. Ceriani, A. Russo, V. Trianni, Top-down and bottom-up approaches to domain knowledge engineering, 2024, <http://dx.doi.org/10.5281/zenodo.14264691>.
- [25] A.G. Nuzzolese, V. Presutti, A. Gangemi, S. Peroni, P. Ciancarini, Aemo: Linked data exploration based on knowledge patterns, *Semant. Web* 8 (1) (2017) 87–112.