

Selected Research Directions in NeuroSymbolic AI

Michael McCain^{1,*}, Antrea Christou¹, Chris Davis Jaldi¹, Elham Ghiasi¹, Susan Shrestha¹, Sumit Sunar¹, Krischal Dhungel¹, Anmol Saini¹, Emily Miller¹, Alexis Ellis¹, Stacie Severyn¹, David Castro¹, Shreyas Casturi¹, Rose Taylor¹, Jonathan Wasky¹ and Cogan Shimizu¹

¹Wright State University, 3640 Colonel Glenn Hwy, Dayton, OH, 45435, USA

Abstract

Large Language Models (LLMs) have demonstrated remarkable capabilities in language understanding and generation, yet persistent challenges in reliability, explainability, and factual grounding limit their adoption in safety-critical and knowledge-intensive domains. This paper presents an examination of how ontologies are being used across three research areas to address the shortcomings of LLMs.

Keywords

Large Language Models (LLMs), Ontology, Knowledge Graphs, Knowledge Representation & Validation, Industry 5.0, Multi-agent system, Neuro-symbolic Artificial Intelligence, Embeddings

1. Introduction

Artificial intelligence, especially LLMs, has shown impressive analytical and pattern-matching capabilities, but reliability, explainability, and trustworthiness remain common deficiencies. These weaknesses become more apparent when domain knowledge is needed to provide accurate, reliable answers. Whether the challenge is the confabulation and unstable reasoning exhibited by LLMs operating without external grounding [1], the difficulty of assigning tasks to heterogeneous human and machine agents in Industry 5.0 workspaces [2, 3], or the broader need for AI and machine-learning systems to reason over coherent, validated factual representations [4, 5, 6], the underlying issues remain the same. Data-driven models, regardless of domain, require a structured symbolic layer that constrains what can be said, what can be done, and how entities relate. That layer is the ontology.

In this work, we present an examination of how ontologies strengthen distinct research areas:

1. Hybrid LLM-KG systems for verifiable reasoning.
2. Ontology-driven multi-agent frameworks for human-machine collaboration in Industry 5.0.
3. Neuro-symbolic AI applications that leverage these structured knowledge representations for domain-specific tasks.

Across all three areas, the introduction of ontologies yields concrete and complementary benefits.

The rest of the paper is structured as follows. Section 2 explores the explainable and verifiable reasoning using ontologies. Section 3 introduces various applications of ontologies in Industry 5.0

Conference here

*Corresponding author.

†These authors contributed equally.

✉ mccain.32@wright.edu (M. McCain); christou.2@wright.edu (A. Christou); jaldi.2@wright.edu (C. D. Jaldi); ghiasi.2@wright.edu (E. Ghiasi); shrestha.167@wright.edu (S. Shrestha); sunar.8@wright.edu (S. Sunar); dhungel.5@wright.edu (K. Dhungel); saini.25@wright.edu (A. Saini); miller.1771@wright.edu (E. Miller); ellis.177@wright.edu (A. Ellis); severyn.5@wright.edu (S. Severyn); castro.31@wright.edu (D. Castro); casturi.2@wright.edu (S. Casturi); taylor.940@wright.edu (R. Taylor); wasky.2@wright.edu (J. Wasky); cogan.shimizu@wright.edu (C. Shimizu)
🆔 0009-0008-3664-6248 (M. McCain); 0009-0003-0428-2439 (A. Christou); 0009-0000-2287-1198 (C. D. Jaldi); 0009-0009-1201-2029 (E. Ghiasi); 0009-0005-0728-6851 (S. Shrestha); 0009-0000-1735-2377 (A. Saini); 0009-0008-4838-9461 (E. Miller); 0009-0000-8098-4262 (A. Ellis); 0009-0008-2705-8946 (S. Severyn); 0009-0006-8593-4188 (D. Castro); 0009-0007-9111-9830 (S. Casturi); 0009-0000-8206-0725 (R. Taylor); 0009-0006-0572-0316 (J. Wasky); 0000-0003-4283-8701 (C. Shimizu)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

settings. Finally, Section 4 discusses the benefits of having structured knowledge integrated into neuro-symbolic systems.

2. LLMs’ Synergy with Knowledge-Grounded, Explainable Tasks

LLMs gained strong attention for their language understanding and generation skills, yet they still produced confabulations or unstable responses when they worked without external knowledge integrations. They further struggle with schema, structural, or symbolic constraints and reasoning, especially when needed to work with explicitly typed rules or relationships typically found in Knowledge Graphs (KGs). Prior work showed that these failures reduced factual grounding, weakened explainability, and limited trust for real-world usecases [1], further making it clear that data-driven models needed stronger support from structured knowledge and clearer reasoning control. Many of these pitfalls generally arose from gaps in how current systems integrated external knowledge. Most pipelines relied on brute-force or unstructured retrieval, which ignored any structural schema constraints. While few systems produced symbolic retrieval plans [7], the lack of validation towards inconsistent structures in parsing pipelines [8] reduced trust in machine-generated graphs. These limitations showed how hybrid LLM systems needed structured representation, schema-aware retrieval, transparent reasoning control, and reliable validation for stable adoption.

In this work, we present a hybrid framework as a position towards more grounded and reliable KG-LLM aware system design, by outlining how (a) controlled knowledge representation (machine and human-readable), (b) structured retrieval, (c) reasoning control through both symbolic and neural constraints, and (d) graph-based validation through reconstruction address the shared problem space across research areas. We summarize how controlled language, ontology structure, and constraints affected ambiguity removal, factual grounding, reasoning behavior, and trust in answer generation by organizing those ideas into a unified view. This work contributes to the central theme: LLMs need richer collaboration with ontologies to support explainable and verifiable reasoning, offering a clear set of directions for building explainable, KG-grounded LLM systems.

2.1. Background and Related Work

Data representations, especially with natural languages, caused various issues with knowledge resolution and factual breakdown. As a result, controlled natural languages (CNLs) were thoroughly motivated to restrict the syntax and semantics of natural language and obtain unambiguous, interpretable texts while preserving accessibility by deterministically translating into discourse structures or further into first-order logic, OWL, or rule languages [9]. This work demonstrated that carefully designed linguistic constraints can resolve semantic ambiguity and enable automated reasoning over the resulting knowledge bases, often critical for mechanisms such as retrieval augmented generation (RAG) systems.

Emerging as a primary strategy for mitigating hallucination through external knowledge integration, recent surveys distinguished naive text-only approaches from more modular architectures that explicitly separate indexing, retrieval, and augmentation stages. Within this space, GraphRAG methods leverage graphs (especially KGs) as structured backends, exploiting their ontology to retrieve subgraphs aligned with the requirements over flat chunks [10] and improve reasoning and traceability compared to naive embedding search [11]. However, most existing pipelines focus on retrieval quality and coverage and pay comparatively little attention to knowledge conceptualization and representation (e.g., reification or schema density) or how minimal, schema-aligned subgraphs should be selected to feed into LLMs.

A large body of work studying multi-hop question answering over KG [12] highlighted the need to combine distributed textual evidence and relational structure while producing explanations [13]. While advanced strategies with LLMs (e.g., chain-of-thought (CoT), in-context learning, graph-of-thought (GoT), etc.) have been shown to elicit stepwise reasoning behavior and improved performance, they were also sensitive to schema exposure, selection, and complexity [14]. These collectively suggested that (i) LLM reasoning quality depends strongly on KG schema and constraint presentation and (ii) prompt design parameters act as powerful but under-optimized levers for controlled symbolic behavior. Beyond

using KGs as cognitive retrieval systems, a complementary line of research uses them as targets or constraints in generation. Such pipelines normalize responses into abstract representations or ontology structures for robustness, logical consistency, discrete factual boundaries, and ontology design pattern alignment [15]. However, most approaches explore them beyond using them as static targets, with limited exploration of graph reconstruction and alignment as a systematic validation step.

2.2. Hybrid Framework and its Components

To bridge the above gaps, we posit a hybrid framework (depicted in Fig. 1), as follows.

Controlled Language for Representing Knowledge (CLaRK) addresses the fundamental problem of semantic ambiguity in knowledge-intensive systems, introduced through implicit, context-dependent, and unrestricted natural language. We adopt the principles of CNLs, specifically Attempto Controlled English (ACE) [9], for syntactically valid and deterministic grounding in first-order logic by formally specifying explicit construction rules for both human-understandable and machine-interpretable expressions. In our work, we utilize LLMs, ACE-compliant prompt instructions, and Attempto Parsing Engine validated feedback loops to break down complex statements, reveal any hidden assumptions, and faithfully preserve the intended domain meaning [16]. These validated representations directly offer better retrieval by aligning content with precise logical structures rather than ambiguous surface text, enabling more reliable knowledge alignment and integration across hybrid reasoning pipelines.

Retrieval for Knowledge Alignment in LLMs (ReKALL) bridges the gap between the unstructured LLM responses and the structured, symbolic requirements of formal ontologies. Specifically, it treats representation-aware, high-precision subgraph retrieval as a first-class design problem rather than a purely engineering detail. Adopting from prior work [11], the workflow embeds domain ontologies, creating a searchable vector space index to identify the key node and perform multi-hop subgraph traversal for relevant schema slices, similar to GoT. This extracted subgraph is responsible for schema-aware querying, ensuring retrieval is logically grounded, and enabling feasibility of higher-order tasks.

Neurosymbolic (NeSy) Reasoning. While CLaRK and ReKALL showed next steps for the knowledge representation and retrieval in itself, a further exploration of reasoning capabilities enabled by LLMs’ neural processing and KGs’ symbolic structuring remained. We designed two components to fill this gap by studying how symbolic scaffolds, schema-only views, few-shot traces, and hyperparameters affect reasoning behavior.

Recent work has shown that the structure and representation of knowledge (either as ontologies [17] or knowledge bases [1]) significantly influence model performance. **Symbology-based Reasoning Effects and Analysis (SyRA)**, modeled this focus even more with factorial experiments and ablation analyses around ontology design aspects such as structure, size, expressivity, coverage, or conceptualization and modeling choices [7]. In contrast, **Neural Reasoning Effects and Adaptation (NeRA)** shifts the focus from how knowledge is modeled to how the model is instructed. While SyRA has prompting and tuning knobs (temperature, few-shot, context-tuning, etc) constant with variations in ontology as an experimental variable, NeRA assumed a stable ontology and explored the effects of LLM tuning parameters. Specifically, we study LLMs’ path-based reasoning ability through SPARQL query generation and path assessment [18]. The central hypothesis of SyRA and NeRA was that the way an ontology is designed and a model is prompted has a strong influence on reasoning over structured knowledge [19]. By treating both variations as first-class experimental factors, they complement




Figure 1: Position LLM-KG hybrid framework portraying (a) CLaRK, (b) ReKALL, (c) SyRA & NeRA and (d) GRAIL.

findings and provide a complete view of the interaction between symbolic knowledge structures and neural models with implications for NeSy systems [7].

Graph Reconstruction and Alignment for Integrity and Validation (GRAIL) further extended this line of work by explicitly framing graph-based reconstruction, alignment, and scoring as mechanisms for detecting hallucination and assessing integrity along a reference KG. In contrast to surface-level verification methods that rely on token overlap or embedding similarity, and adopting from recent works [8, 20], an emphasis here was placed on structural and relational fidelity, enabling more fine-grained validation. This approach is particularly effective for complex reasoning tasks, multi-hop factual queries, and explanations in which correctness depends on the coherence of relationships rather than isolated facts and constitutes a promising direction for improving trust, interpretability, and factual integrity in LLM-generated responses.

2.3. Opportunities and Limitations

The proposed framework offers several opportunities. By cross-integrating with each other, it provides a unified lens for LLM interaction with symbolic structures and enables systematic interaction studies that prior work has only examined in isolation. For example, ACE-style constraint propagation into retrieval precision, representational choice influencing downstream GraphRAG and reasoning performance, or even SyRA/NeRA reasoning with graph reconstructed validation draws evidence for more robustness, integrity, and trustworthiness, an opportunity largely unaddressed in existing hybrid systems.

Despite these benefits, it also faces limitations. CNL imposes authoring overhead and may require very delicate domain adoption. High precision retrievals and schema-based prompting heavily depend on KG’s quality and structural information availability, which varies in representation across domains. Furthermore, LLM inference remains sensitive to architectural versions and tuning parameters, complicating reproducibility and general adaptability. Finally, graph reconstruction provides a validation path, but evaluating entire semantics, especially at scale, remains non-trivial.

Overall, these next steps, opportunities, and limitations illustrate both the need and the potential for systematic hybrid approaches to LLM-KG integration. Our framework serves as a foundation for such work and provides a structured basis for future empirical evaluation.

3. Hybrid Artificial Intelligence Agents

3.1. Introduction

Industry 5.0 is the natural next step for the integration of the technological achievements made with Industry 4.0 into shared workspaces with humans [2]. The core of Industry 5.0 is human-centered, highlighting how these new technologies can work with humans, bringing humans back into the loop to increase performance, safety, and fairness in the workplace while also minimizing the cost and time of production. Considering the human-centric nature of Industry 5.0, there are interesting intersections within domains like explainable artificial intelligence (XAI), human factors principles, cognitive sciences, and collaborative robotics (co-bots). Specifically, we aim to investigate how we can organize agents for effective and efficient collaboration while also taking into consideration how agents can interact with each other seamlessly. One approach is *the multi-objective task allocation*, which involves the ability to assign tasks to other agents (human or machine) with consideration of the agents capabilities to perform the assigned task [3]. To effectively allocate tasks to different agents, it is important to have a hierarchical chain of command that consists of AI capabilities and human-in-the-loop approaches, an accessible knowledge base of agent capabilities, and the ability for agents to interact and communicate with one another. In this paper, we combine the research efforts of 4 ongoing projects to create a framework that can effectively allocate multi-objective tasks to agents, facilitate seamless collaborative interactions between agents, and create accessible knowledge of agent capabilities, all of which, when combined, will reflect the initiatives of Industry 5.0.

3.2. TASK-MATS: Task Adaptation with Shared Knowledge for Multi-Agent Teaming Systems

Industries such as manufacturing, which rely heavily on robotic systems, have expressed a desire for more complex autonomous systems capable of human-robot collaboration and reasoning [21, 22]. To fulfill this desire, the Knowledge and Semantic Technologies (KASTLE) Laboratory¹ is working toward the development of a framework that enables an AI Overseer (AIO) to orchestrate multi-agent teams (MATS) comprising autonomous robots, digital systems, and other agents to complete tasks in unfamiliar settings without the need for additional training based on previous experiences or data from a continuously updated knowledge base. This framework has been labeled Task Adaptation with Shared Knowledge for Multi-Agent Teaming Systems (TASK-MATS).² To enable this orchestration, the proposed framework would make use of symbolic systems, such as an ontology, to reason about the world in which the orchestrator and agents operate. As new knowledge is garnered through practical observation, inferences and inductions can be made [23]. From this, the AIO will be able to reason about the agents available to it, their capabilities, and how to effectively utilize them to achieve long-term goals and immediate objectives. In addition, the AIO will begin to learn the full capabilities of each MAS, enabling a process for approaching assigned tasks at different levels. Although the current work focuses specifically on the domain of manufacturing [24], the architecture and design of the TASK-MATS framework would facilitate cross-domain cooperation or specialization for a single knowledge realm, depending on user needs.

3.3. KGWRAPS: A Knowledge Graph-powered Research Assistant for Polymer Science

The accelerated development of autonomous systems possesses the potential to greatly advance various scientific fields. In polymer science, however, their integration remains limited, with repetitive manual experimentation and data curation prevailing [25]. By introducing human biases and error, these aspects diminish the replicability of existing polymers, while their time-consuming nature impedes the rate of discovery of new polymers, ultimately stagnating the field.

Thus, we introduce a Knowledge Graph-poWered Research Assistant for Polymer Science (KGWRAPS),³ a neurosymbolic artificial intelligence (NeSy AI) agent outfitted with large language models (LLMs) and knowledge graphs (KGs) to support domain experts in the construction and execution of polymer experiments. By combining conversational capabilities [26] with complex domain knowledge [27], this research assistant will expedite work in polymer science.

In the trend toward Industry 5.0, the involvement of domain experts in automated experimental development becomes paramount, not only to verify the accuracy of an autonomous agent's experimental processes but also to ensure their safety, especially as the limitations posed by AI guardrails manifest more evidently [28]. Thus, the integration of a robot powered by KGWRAPS, along with a human expert, into the TASK-MATS framework holds significant promise in leveraging both the accuracy, precision, and speed of modern technology and the experience and human awareness of domain experts to advance the field both quickly and safely.

3.4. A Knowledge Graph for the Detection of Malware Behavior in Executables

The process of software reverse engineering involves decompiling and disassembling a program to get the original source code and is a useful method for detecting malware in software. Leveraging reverse engineering software like *Ghidra* or *IDA pro* allows a person to look through the decompilation and disassembly of a file to identify the functions and variables within the software, providing an overall summary of functionalities [29]. While there already exist tools to analyze files in this way, it is often a tedious and time-consuming process. The process is further complicated by some programs containing

¹<https://kastle-lab.github.io>

²<https://kastle-lab.github.io/TASK-MATS-web>

³<https://kastle-lab.github.io/kgwraps-web>

data obfuscation, concealing data to hide their purpose. Therefore, in this project, we introduce a KG⁴ based on the symbol table in Ghidra to help simplify the process of detecting malicious code.

The KG's integration into TASK-MATS will assist human agents in the field of malware analysis and reverse software engineering, furthering Industry 5.0. This project can be leveraged by finding patterns at the disassembly level, as well as other data connections that can be found through defined strings and other parts of the Ghidra analysis. Specific patterns of interest include the frequency of certain kinds of instructions, like jump or halt instructions, cases of instructions being used incorrectly, and calls to manipulate file systems or other data [30]. Other things to look for include seeing what imports the file uses, how much of the data is obfuscated, and if it connects to the network in any way. These patterns can be used to detect malicious behavior in executable files to facilitate identifying malware. Utilizing collaboration between humans and machines in this way will help increase workplace safety and device security going into Industry 5.0.

3.5. OntoPret: An Ontology for the Interpretation of Human Behavior

A critical need for the future of hybrid machine teaming is the machines' ability to interpret the complex behaviors of their human counterparts. Bridging the gap between techno-centric robotic frameworks and descriptive behavioral ontologies becomes an important step for the continued push towards Industry 5.0. We introduce our project *OntoPret: An Ontology for the Interpretation of Human Behavior* [31]. OntoPret⁵ is grounded in the cognitive sciences and modular engineering methodology, providing a formal, machine-processable framework that is aimed at classifying behaviors like task deviation and deceptive actions. OntoPret was created using the Modular Ontology Modeling methodology (MOMo) [32], making for seamless integration and re-usability of other ontology frameworks. Therefore, when integrated into other, more techno-centric frameworks like TASK-MATS, OntoPret provides the human-centric values of Industry 5.0. Currently, OntoPret highlights behaviors associated with deviation and deception, but as a result of the methodology used to develop OntoPret, more behaviors can be added, creating a larger knowledge repository for the machines to query regarding their human counterparts.

3.6. Conclusions

In this work, we have proposed the potential integration of ontology-driven projects across various domains to establish a modular framework usable in Industry 5.0. With TASK-MATS more broadly framing agent interactions through an AIO, specific AI agent projects (such as KGWRAPS) can be included to support refined conversational models that guide the domain-specific aspects of these interactions in collaborative settings. However, the potential for integration surpasses simply interactive considerations. In the context of domains requiring digital agents, such as those in fields involving malware behavior analysis, such a modular framework could be enhanced to facilitate these use cases and assist humans in identifying malicious software. Furthermore, via integration with other more humanistic, ontology-driven works such as OntoPret, the TASK-MATS framework could be better enabled to accommodate environments where safe and effective human-robot collaborations are essential. Thus, through a multidisciplinary analysis of the advancement of multi-agent teams, the trend toward Industry 5.0 substantiates the promise of significant progress.

4. Introduction

Artificial intelligence (AI) and machine learning (ML) applications, depend on structured data for reasoning, inference, and decision-making. Knowledge Graphs (KGs), which represent things and their relationships in a graph-based manner, can assist in maintaining these properties [4].

Structured knowledge has become extremely important with the rapid emergence of AI models that rely on enormous volumes of data to learn, reason, and make judgments. The reliability of the

⁴<https://github.com/22lavonne/Malware-Behavior-Knowledge-Graph>

⁵<https://github.com/kastle-lab/OntoPret-Ontology-for-the-Interpretation-of-Human-Behavior>

models is greatly affected by the representation, organization, and validity of facts as these models become stronger over time. This means that, in addition to using organized information, it is crucial to carefully consider what constitutes “good” or “effective” architecture while teaching and assessing them. Well-designed data help reduce ambiguity, increase interpretability, and allow more meaningful evaluation of model behavior, especially in tasks that depend on contextual knowledge rather than simple pattern matching.

The significance of structured information in improving deep learning understanding has been identified in recent years by Neuro-Symbolic AI, a combination of neural networks and symbolic reasoning [5, 6].

4.1. Background

4.1.1. Neural Models

Neural models are computer models used in machine learning that are inspired by how the human brain is “wired” [33]. They are made up of many units called neurons that are connected to each other and organized in layers. These neurons take input data, process it using weights and mathematical functions, and pass the results forward to produce an output. Neural networks are especially good at learning complex patterns in data, which is why they are commonly used for tasks like image recognition, speech recognition, and prediction problems. They are a fundamental component of supervised, unsupervised, and deep learning approaches in machine learning [34].

4.1.2. Symbolic Reasoning

The integrated aspect of symbolic reasoning suggests that reasoning is shaped not just by abstract rules, but by how external symbols and structures are observed and interacted with [35]. Applied to knowledge graphs, this implies that reasoning over entities and relationships is influenced by the graph’s structure and organization, with the layout and connections of nodes supporting inference and shortcuts. In this way, symbolic reasoning in knowledge graphs combines logical operations with the structural and observable cues inherent in the graph.

4.1.3. Artificial Intelligence


Focuses on creating networks capable of interpreting, organizing, and reasoning over structured knowledge to perform tasks that typically require human intelligence [36]. These tasks include learning from data, inferring new relationships, answering complex queries, and integrating heterogeneous information. AI encompasses approaches ranging from symbolic reasoning over ontologies to machine learning techniques for embeddings and machine learning tasks, enabling adaptations and predictions over complex data. Its development draws on multiple disciplines, including graph theory, cognitive science, and data engineering, and supports applications such as semantic search, decision support, and knowledge management.

4.2. Applications

4.3. Pattern Breakdown for Early Detection of Harmful Algal Blooms

Monitoring harmful algal bloom has traditionally relied on individual indicators of water-quality evaluated against fixed thresholds [37].

In many lakes and reservoirs, water-quality variables tend to move together in stable patterns of multivariate coupling, and growing evidence suggests that these relationships begin to break down before a bloom becomes



visible. Tracking when and how this coupling changes offers a way to recognize early warning signs.

Future work should explore how well these patterns transfer between similar sites and how such signals can be reliably used for real-time monitoring.

4.3.1. Graph Structure and Knowledge Graph Embeddings

Controlled experiments are constructed using both synthetic and real-world knowledge graphs to systematically evaluate the behavior of knowledge graph embeddings. By analyzing model performance across graph properties that are fundamentally symbolic; such as graph size, degree centrality, connectivity, and relational structure, the results demonstrate that embedding behavior is strongly influenced by the structure of the underlying knowledge representation rather than by model architecture alone [38].

To further isolate and analyze the effects of graph structure and known symbolic structure, synthetic knowledge graphs were constructed by assigning each node a random real value $x_i \in [0, 1]$ and deriving symbolic structure through pairwise *LessThan* relations and recursively defined *PartOf* window relations, after which the numeric values are discarded while preserving the induced structure. These graphs are embedded and analyzed via principal component analysis [39] and a multilayer perceptron [40] applied to the first two principal components.

4.3.2. Evaluating LLM reasoning with registered-based prompting

Determining whether LLMs can plausibly reason or not is complicated by the fact that researchers have disparate tasks to measure reasoning [41].

Therefore, time was spent determining a series of tasks that could be fashioned into a test suite for an LLM reasoning benchmark dataset, as well as what metrics this dataset would require. For example, some of the tasks involved would be rewriting logical formulas with variable substitution, or modeling arithmetic problems which require multiple steps or stages of computation.

What would happen when an LLM is prompted to use a domain-specific language (DSL) that involves the use of "registers", which can hold intermediate values or computations, simulating a CPU's memory?

More work needs to be done on standardization of tasks/benchmark and finalizing the prompt strategies before testing an LLM with it.

Future work could include having an LLM use a KG with examples of solved problems using the DSL.

4.3.3. LLM assisted sentiment analysis

In the world of AI, sentiment analysis has been treated as a giant statistical guessing game [42]. Often feeding millions of tweets into a "black box" until it understands that "terrible" is negative. Large language models often struggle with sarcasm, tone, and certain phrases. To better fine tune the LLM, it is fed the sentiment results that a group of human researchers came up with for the same dataset. With this fine-tuning in mind, the set of reviews are rerun through the LLM and the output is compared against the group of researchers. This creates a recursive process, looping until both the LLM and human researchers align in sentiment. This convergence between LLMs and human researchers does more than just correct errors, it helps the LLM develop internal heuristics that mimic human reasoning.

Consequently, the LLM stops counting the sentiment of each individual word and becomes a reliable interpreter of nuanced language.

4.4. Limitations & Future Work

Limitations exist in all applications discussed above in terms of static workflows, ad hoc parameter tuning, and the need for task-specific adjustments. Future developments aim to create dynamic experimental pipelines that allow easy fine-tuning of parameters for pattern detection, prompt engineering and kge evaluation techniques.

5. Acknowledgments

Acknowledgments

Chris Davis Jaldi acknowledges support from DHS award FA8075-18-D-0002/FA8075-21-F-0074. Elham Ghiasi and Cogan Shimizu acknowledge support from NSF award #2333532. Icons are from Flaticon. Michael McCain and Cogan Shimizu acknowledge support from DAGSI RX26-24: “TASK-MATS: Task Adaptation with Shared Knowledge for Multi-Agent Teaming Systems.” Anmol Saini and Cogan Shimizu acknowledge support from DAGSI RX24-22: “KGWRAPS: A Knowledge Graph-Powered Research Assistant for Polymer Science.” Alexis Ellis and Cogan Shimizu acknowledge support from the National Science Foundation under award #2333532 “Proto-OKN Theme 3: An Education Gateway for the Proto-OKN (EduGate).” Antrea Christou, David Castro, Jonathan Wasky and Cogan Shimizu acknowledge support from the National Science Foundation (NSF) under Grant #2333532; Proto-OKN Theme 3: An Education Gateway for the Proto-OKN.

Declaration on Generative AI

During the preparation of this work, the authors used Grammarly for grammar and spell checks.

References

- [1] J. Sequeda, D. Allemang, B. Jacob, Knowledge graphs as a source of trust for llm-powered enterprise question answering, *Journal of Web Semantics* 85 (2025) 100858.
- [2] M. T. Islam, K. Sepanloo, S. Woo, S. H. Woo, Y.-J. Son, A review of the industry 4.0 to 5.0 transition: exploring the intersection, challenges, and opportunities of technology and human-machine collaboration, *Machines* 13 (2025) 267.
- [3] M. Calzavara, M. Faccio, I. Granata, Multi-objective task allocation for collaborative robot systems with an industry 5.0 human-centered perspective, *The International Journal of Advanced Manufacturing Technology* 128 (2023) 297–314.
- [4] T. R. Gruber, A translation approach to portable ontology specifications, *Knowledge acquisition* 5 (1993) 199–220.
- [5] P. Hitzler, M. K. Sarker (Eds.), *Neuro-Symbolic Artificial Intelligence: The State of the Art*, volume 342 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2021. URL: <https://doi.org/10.3233/FAIA342>. doi:10.3233/FAIA342.
- [6] A. Sheth, K. Roy, M. Gaur, Neurosymbolic artificial intelligence (why, what, and how), *IEEE Intelligent Systems* 38 (2023) 56–62. doi:10.1109/MIS.2023.3268724.
- [7] C. D. Jaldi, et al., Knowledge conceptualization impacts rag efficacy, in: *International Knowledge Graph and Semantic Web Conference*, Springer, 2025, pp. 208–223.
- [8] A. Gangemi, A. Graciotti, A. Meloni, et al., Text2amr2fred, converting text into rdf/owl knowledge graphs via abstract meaning representation, *Knowledge and Information Systems* 68 (2026) 47.

- [9] N. E. Fuchs, K. Kaljurand, T. Kuhn, *Attempto controlled english for knowledge representation*, in: *Reasoning Web: 4th International Summer School 2008, Venice, Italy, September 7-11, 2008, Tutorial Lectures*, Springer, 2008, pp. 104–124.
- [10] Q. Zhang, S. Chen, Y. Bei, et al., *A survey of graph retrieval-augmented generation for customized large language models*, *arXiv preprint arXiv:2501.13958* (2025).
- [11] N. Matsumoto, J. Moran, H. Choi, M. E. Hernandez, et al., *Kragen: a knowledge graph-enhanced rag framework for biomedical problem solving using large language models*, *Bioinformatics* 40 (2024) btac353. doi:10.1093/bioinformatics/btac353.
- [12] A. Saxena, A. Tripathi, P. Talukdar, *Improving multi-hop question answering over knowledge graphs using knowledge base embeddings*, in: *Proceedings of the 58th annual meeting of the association for computational linguistics*, 2020, pp. 4498–4507.
- [13] Z. Yang, P. Qi, S. Zhang, Y. Bengio, W. Cohen, R. Salakhutdinov, C. D. Manning, *Hotpotqa: A dataset for diverse, explainable multi-hop question answering*, in: *Proceedings of the 2018 conference on empirical methods in natural language processing*, 2018, pp. 2369–2380.
- [14] A. Gashkov, A. Perevalov, M. Eltsova, A. Both, *Sparql query generation with llms: Measuring the impact of training data memorization and knowledge injection*, in: *International Conference on Web Engineering*, Springer, 2025, pp. 177–192.
- [15] C. Shimizu, P. Hitzler, *Accelerating knowledge graph and ontology engineering with large language models*, *Journal of Web Semantics* 85 (2025) 100862.
- [16] G. Garzo, A. Palumbo, *Human-in-the-loop: Legal knowledge formalization in attempto controlled english*, in: *2025 13th International Symposium on Digital Forensics and Security (ISDFS)*, IEEE, 2025, pp. 1–6.
- [17] A. Christou, C. Shimizu, *Experiments in graph structure and knowledge graph embeddings*, *Neurosymbolic Artif. Intell* (2026).
- [18] C. D. Jaldi, *Impact of Graph Structures for RAG Outcomes in LLMs*, Master's thesis, Wright State University, 2025.
- [19] H. Tian, C. Wang, B. Yang, L. Zhang, Y. Liu, *A taxonomy of prompt defects in llm systems*, *arXiv preprint arXiv:2509.14404* (2025).
- [20] A. Gangemi, A. G. Nuzzolese, *Logic augmented generation*, *Journal of Web Semantics* (2025). doi:10.1016/J.WEBSEM.2024.100859.
- [21] H. Du, S. Thudumu, R. Vasa, K. Mouzakis, *A survey on context-aware multi-agent systems: Techniques, challenges and future directions*, 2024. URL: <https://arxiv.org/abs/2402.01968v2>.
- [22] E. Aguado, V. Gomez, M. Hernando, C. Rossi, R. Sanz, *A survey of ontology-enabled processes for dependable robot autonomy*, *Frontiers in Robotics and AI* 11 (2024). doi:10.3389/frobt.2024.1377897.
- [23] L. Bühmann, J. Lehmann, P. Westphal, *DI-learner—a framework for inductive learning on the semantic web*, *Journal of Web Semantics* 39 (2016) 15–24. URL: <https://www.sciencedirect.com/science/article/pii/S157082681630018X>. doi:<https://doi.org/10.1016/j.websem.2016.06.001>.
- [24] M. McCain, C. D. Jaldi, S. Shrestha, S. Casturi, C. Shimizu, *Towards a modular ontology for autonomous robotic orchestration*, volume In-Print, *International Workshop on Ontologies for Autonomous Robotics 2025, The International Association for Ontology and its Applications (IAOA)*, 2025.
- [25] T. Schuett, P. Endres, T. Standau, S. Zechel, R. Q. Albuquerque, C. Brütting, H. Ruckdäschel, U. S. Schubert, *Application of Digital Methods in Polymer Science and Engineering*, *Advanced Functional Materials* 34 (2024) 2309844. doi:10.1002/adfm.202309844.
- [26] A. Saini, J. G. Ethier, C. Shimizu, *An ontology for conversations with virtual research assistants*, in: *2024 IEEE 36th International Conference on Tools with Artificial Intelligence (ICTAI)*, 2024, pp. 181–186. doi:10.1109/ICTAI62512.2024.00034.
- [27] A. Saini, C. D. Jaldi, J. Ethier, C. Shimizu, *Research directions for ontology-guided domain-specific knowledge graph population using llms*, in: B. Villazón-Terrazas, F. Ortiz-Rodriguez, S. Tiwari, T. Riechert, E. Marx (Eds.), *Knowledge Graphs and Semantic Web*, Springer Nature Switzerland, Cham, 2026, pp. 137–145.
- [28] S. A. Akheel, *Guardrails for large language models: A review of techniques and challenges*, *J Artif*

Intell Mach Learn & Data Sci 3 (2025) 2504–2512.

- [29] R. Rohleder, *Hands-on ghidra - a tutorial about the software reverse engineering framework*, in: *Proceedings of the 3rd ACM Workshop on Software Protection, SPRO'19*, Association for Computing Machinery, New York, NY, USA, 2019, p. 77–78. URL: <https://doi.org/10.1145/3338503.3357725>. doi:10.1145/3338503.3357725.
- [30] S. Severyn, *Adapting linguistic deception cues for malware detection*, 2014. URL: <https://api.semanticscholar.org/CorpusID:56043342>.
- [31] A. Ellis, S. Severyn, F. Novakazi, H. Banaee, C. Shimizu, *Ontopret: An ontology for the interpretation of human behavior*, *arXiv preprint arXiv:2510.23553* (2025).
- [32] C. Shimizu, K. Hammar, P. Hitzler, *Modular ontology modeling*, *Semantic Web* 14 (2023) 459–489.
- [33] R. Qamar, B. A. Zardari, *Artificial neural networks: An overview*, *Mesopotamian Journal of Computer Science* 2023 (2023) 124–133. URL: <https://mesopotamian.press/journals/index.php/cs/article/view/118>. doi:10.58496/MJCSC/2023/015.
- [34] I. El Naqa, M. J. Murphy, *What Is Machine Learning?*, Springer International Publishing, Cham, 2015, pp. 3–11. URL: https://doi.org/10.1007/978-3-319-18305-3_1. doi:10.1007/978-3-319-18305-3_1.
- [35] D. Landy, C. Allen, C. Zednik, *A perceptual account of symbolic reasoning*, *Frontiers in Psychology* 5 (2014) 275. doi:10.3389/fpsyg.2014.00275.
- [36] V. Yashchenko, *Artificial intelligence theory (basic concepts)*, in: *2014 Science and Information Conference*, 2014, pp. 473–480. doi:10.1109/SAI.2014.6918230.
- [37] S. Xu, Z. Ye, S.-C. Hsu, X. Liu, C. Zheng, *Development of data-driven algal bloom alert models with low temporal resolution data and application to hong kong rivers*, *Journal of Hydrology: Regional Studies* 57 (2025) 102108. URL: <https://www.sciencedirect.com/science/article/pii/S2214581824004579>. doi:<https://doi.org/10.1016/j.ejrh.2024.102108>.
- [38] A. Christou, C. Shimizu, *Experiments in graph structure and knowledge graph embeddings*, *Neurosymbolic Artificial Intelligence* (2024). In press.
- [39] F. Kherif, A. Latypova, *Chapter 12 - principal component analysis*, in: A. Mechelli, S. Vieira (Eds.), *Machine Learning*, Academic Press, 2020, pp. 209–225. URL: <https://www.sciencedirect.com/science/article/pii/B9780128157398000122>. doi:<https://doi.org/10.1016/B978-0-12-815739-8.00012-2>.
- [40] J. Singh, R. Banerjee, *A study on single and multi-layer perceptron neural network*, in: *2019 3rd International Conference on Computing Methodologies and Communication (ICCMC)*, 2019, pp. 35–40. doi:10.1109/ICCMC.2019.8819775.
- [41] P. Shojae*, I. Mirzadeh*, K. Alizadeh, M. Horton, S. Bengio, M. Farajtabar, *The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity*, in: *NeurIPS*, 2025. URL: <https://arxiv.org/abs/2506.06941>.
- [42] A. Areshey, H. Mathkour, *Exploring transformer models for sentiment classification: A comparison of bert, roberta, albert, distilbert, and xlnet*, *Expert Systems* 41 (2024). doi:10.1111/exsy.13701.

A. Online Resources

- *Sentiment Analysis*
- *KGE Experiments*