

# Thesis Proposal: A Multi-Agent System for Ontology-Based Perspective-Aware Knowledge Extraction

**Luiz do Valle Miranda**

Department of Human-Centered AI  
Institute of Applied Computer Science  
Jagiellonian University  
Kraków, Poland  
luiz.miranda@uj.edu.pl

**Grzegorz J. Nalepa**

Department of Human-Centered AI  
Institute of Applied Computer Science  
Jagiellonian University  
Kraków, Poland  
grzegorz.j.nalepa@uj.edu.pl

## Abstract

This thesis investigates how polyvocal ontologies and Large Language Model (LLM) based Multi-Agent Systems (MAS) can operationalize perspective-aware knowledge extraction, preserving conflicting stakeholder interpretations as epistemically separable, queryable Knowledge Graphs (KGs). Current AI systems consolidate multiple perspectives into singular, decontextualized schemas, introducing representational bias and information loss. We propose a systematic framework addressing three interconnected research questions: (1) how to generate polyvocal ontology design patterns for high-stakes domains; (2) how to architect LLM-based MAS that extract perspective-conditioned facts while maintaining schema coherence and provenance traceability; and (3) whether such extractions achieve semantic diversity without sacrificing KG integrity. Evaluation is proposed on medical datasets, conducted with domain experts, to demonstrate the feasibility of perspective-aware extraction as a principled alternative to consensus-oriented KGs. Expected contributions include polyvocal ontology patterns, an ontology-orchestrated MAS extraction framework with auditable provenance, and empirical validation.

## 1 Introduction

Most contemporary AI systems have a fundamental starting point: the consolidation of multiple sources, viewpoints, and framings into a unified perspective. On the one hand, knowledge engineers design schemas to prevent contradiction. On the other hand, LLM training harnesses human feedback to align toward particular truth standards.

This framework presumes that knowledge should be singular, decontextualized, and free from perspective carries profound consequences. It leads to representational bias, where dominant perspectives in training data override marginalized viewpoints, and the resulting artifacts propagate these

biases at scale (Gallegos et al., 2024). It introduces information loss through decontextualization, thus removing the source, stakeholder, and interpretive context that situated knowledge requires (Kraft and Soulier, 2024). It prevents fine-grained system modeling, as legitimate domain variations cannot be represented and queried separately (van Erp and de Boer, 2021). Finally, it limits downstream applicability—systems built on decontextualized knowledge cannot easily adapt to changing stakeholder requirements (Kay et al., 2025).

Recent work has begun to articulate the value of preserving and processing multiple, sometimes contradictory perspectives as structurally co-valid in an AI system. This so-called perspectivist turn (Cabitza et al., 2023) has created trends in different subfields of AI research. In Natural Language Processing (NLP), perspectivism has emerged as a paradigm that rejects label aggregation in favor of preserving individual annotator judgments, particularly for subjective tasks like hate speech detection, sentiment analysis, and emotion recognition. However, perspectivism in NLP has remained largely confined to annotation and evaluation methodologies, without propagating downstream to knowledge-based systems (Frenda et al., 2024).

In KGs, there is increasing interest in so-called polyvocality in ontologies, that is, the explicit representation and integration of multiple, sometimes conflicting perspectives, voices, or standpoints within a single KG (van Erp and de Boer, 2021; Shoilee et al., 2023). However, existing polyvocal KGs remain largely confined to cultural heritage, with limited transferability to other domains, such as open science, healthcare, policy, or manufacturing.

In LLM research, such a turn has found conceptual grounding in the definition by Kovač et al. of LLMs as a “superposition of perspectives” (Kovač et al., 2023). Such framing has sparked a

perspective-aware methods designed to condition retrieval rankings and LLM outputs on user standpoints (Hayati et al., 2024). Furthermore, MAS architecture have been used for coordination of LLM assessment of different perspectives (Saadaoui and Alonso, 2025; Park et al., 2025). Yet perspective creation lacks mechanisms based on explicitly defined knowledge, thus potentially falling into unqueryable stereotypes.

To address the main challenge of creating perspective-aware AI systems, and the specific challenges of the scarcity of polyvocal resources for knowledge engineering, and principled control over perspectivist knowledge extraction, I propose exploring the prospects of cross-fertilization between developments in polyvocal KGs and LLMs. The remainder of this thesis proposal has the following structure: Section 2 further explores related works on perspectivist AI. Section 3 details the research questions and expected contributions. Section 4 presents some challenges to this work. Section 5 presents a conclusion for the proposal.

## 2 Related Works

### 2.1 Ontology Generation, LLMs and Polyvocality

A rising trend at the intersection of LLMs and Semantic Web (SW) technologies is the use of the former in the process of ontology engineering. Recent work demonstrates that LLMs can effectively assist in multiple ontology engineering tasks, from generating OWL ontology drafts from requirements (via competency questions and user stories) to extracting taxonomies and semantic relations from unstructured text (Lippolis et al., 2025). Furthermore, Zhang et al. (2025) present a framework for conversational ontology engineering towards the collection of requirements from the perspective of different stakeholders and domain experts.

Shimizu and Hitzler (2025) present arguments that modular approaches to ontology design amplify these LLM advantages. First, LLMs can potentially generate high-quality Ontology Design Patterns (ODPs). Second, using previously designed patterns and a module-by-module approach makes the instructions to LLM easier to follow, thereby potentially increasing the quality and consistency of generated ontologies.

While there are clear indications of the potential of using LLMs in the ontology engineering process, there are significant challenges for the case

of polyvocality. First, there is a scarcity of ODPs for polyvocality to be potentially used in prompting LLMs for the generation of polyvocal ontology modules. The exception being (Gangemi and Presutti, 2022), who presents a universal *Cognitive Perspectivization* ontology module that operationalizes perspectives as first-class, reusable ontology design constructs. Second, as mentioned above, there are few available polyvocal solutions for ontologies besides the field of cultural heritage, limiting the use of existing solutions for guiding LLMs.

### 2.2 Perspective-Aware Knowledge Extraction

Another locus of interaction between KGs and LLMs is in the field of knowledge extraction. One key approach is based on ontology-guided prompting, which steers extraction toward schema-consistent outputs. Khorshidi et al. (2025) present a system that dynamically generates ontology snippets tailored to each entity type, aligning extractions with schema constraints and enabling scalable, type-consistent fact extraction across 195 predicates. The system supports both batch and streaming modes, processing over 9 million Wikipedia pages and ingesting 19 million high-confidence facts with 98.8% precision.

A complementary approach for knowledge extraction leveraging modular ontology guidance is presented by Norouzi et al. (2024). Using a three-stage pipeline—module-guided summarization, retrieval-augmented generation, and few-shot prompting—LLMs achieved approximately 90% triple extraction coverage across different prompting strategies and ontology schemas. This demonstrates that the presence of modular ontological structure in knowledge extraction pipelines significantly improves quality.

While these systems have shown substantial potential for schema-consistent knowledge extraction and scalable fact population across diverse domains, there has been no systematic framework yet for perspective-aware KG population, that is, the structured extraction of multi-perspective knowledge using LLMs conditioned on perspective-specific ontological schemas.

### 2.3 Ontologies and LLM-based MAS

Finally, the third locus of interaction between ontologies and language models lies in multi-agent systems. LLMs have been increasingly used as the backbone for goal-oriented agents, that is, autonomous computational entities that perceive local

information, make decisions, and interact via decentralized coordination, communication, or competition in shared environments.

The integration of KGs and perspectives for MAS has followed two separate paths. On the one hand, there has been developments in LLM-based MAS explicitly incorporating multi-perspective assessment to enhance reasoning depth and collective intelligence. CIR3 (Collective Intentional Reading through Reflection and Refinement) presented by [Saadaoui and Alonso \(2025\)](#) demonstrates this through coordinated agents assigned distinct perspective-specific roles based on identified document subtopics. Rather than using a single agent or homogeneous agent teams, CIR3 dynamically identifies conceptually coherent perspectives within a document—e.g., "financial vs. regulatory interpretation" or "clinical vs. epidemiological viewpoint"—and assigns specialized writer agents to each perspective.

On the other hand, KGs have been used for coordinating multi-agent systems across scientific discovery and clinical diagnosis. In science, SciAgents coordinates LLM agents over large-scale KGs to autonomously generate and refine hypotheses; the graph acts as a structured prior that organizes exploration and justifies inferences through explicit paths between concepts and relations ([Ghifarollahi and Buehler, 2025](#)). In healthcare, KG4Diagnosis adopts a hierarchical multi-agent architecture—general practitioner triage plus specialist agents—coordinated via a clinical KG that guides discipline selection, constrains diagnostic reasoning, and preserves verifiable evidence paths for recommendations ([Zuo et al., 2024](#)).

However, no current system has integrated these two advances: using perspective-specific ontological modules to ground multi-agent extraction, such that different agent extract different yet schema-consistent facts depending on their assigned perspective, and those extractions are maintained as epistemically separable knowledge within a polyvocal knowledge infrastructure.

### 3 Research Questions

The main goal of this research is to develop and evaluate a framework for perspective-aware knowledge extraction that uses polyvocal ontologies and LLM-based MAS to populate KGs while preserving conflicting stakeholder perspectives as epistemically separable, queryable representations.

The research goal is decomposed into the following interconnected research questions:

- RQ1: How can polyvocal ontology structures enable perspective-conditioned knowledge extraction?
- RQ2: How can LLM-based MAS leverage polyvocal ontologies to extract schema-consistent, perspective-specific facts?
- RQ3: Do perspective-conditioned extractions achieve semantic diversity while maintaining schema coherence and inter-agent consistency?

The main field of application for the proposed framework is medical data, where diverse stakeholder perspectives (e.g., clinical vs. epidemiological; different specializations) often conflict yet require interoperable representation. The main dataset for this research is processed in collaboration with the Jagiellonian University Medical Faculty under strict ethical guidelines, including anonymization, informed consent where applicable, and compliance with data protection regulations. Additionally, publicly available datasets will be used for increased reproducibility<sup>1</sup>.

In the rest of this section we elaborate on each research question by presenting its constituent sub-questions, methodological approach, evaluation framework, and anticipated contributions.

#### 3.1 RQ1: Polyvocal Ontology Design

As a pre-requisite for the MAS-based perspective-aware extraction (RQ2), we must establish ontological structures that represent and operationalize perspectives in machine-readable, reusable form. This addresses a critical gap: while polyvocal KGs exist primarily in cultural heritage ([van Erp and de Boer, 2021](#)), transferable ODPs for other domains remain scarce<sup>2</sup>.

RQ1 decomposes into two sub-questions:

- **RQ1.1:** How can LLMs generate domain-specific polyvocal ODPs capturing relevant perspective distinctions?

<sup>1</sup>At the moment, we are evaluating the applicability of MIMIC-III ([Johnson et al., 2016](#)) for our research case

<sup>2</sup>A source of inspiration for the development of such ODPs can be both "qualifiers" from Wikidata and the currently diverse representation over multilingual ontologies in DBpedia.

- **RQ1.2:** How can polyvocal ODPs and LLMs—as a “superposition of perspectives”—retrofit existing univocal ontologies (e.g., SNOMED CT) with perspectivization modules?

**Methodology:** We develop LLM-driven pipelines for (1) novel ODP generation from domain requirements (competency questions, stakeholder interviews) and (2) *retrofitting* existing ontologies via perspective-conditioned prompting. Modular ontology design (Shimizu and Hitzler, 2025) guides the process, ensuring shallow semantics and prompt consistency. Furthermore, focusing on human-defined patterns, we can mitigate the cost of re-generating whole ontologies when domain knowledge changes or evolves.

**Evaluation:** Domain experts from medicine assess generated ODPs against perspective-specific competency questions (e.g., “What clinical vs. epidemiological interpretations of symptom-diagnosis relationships exist?”). Success metrics include:

- **Completeness:** Coverage of stakeholder-relevant perspective distinctions
- **Reusability:** Applicability of ODPs across multiple medical scenarios
- **Expert validation:** Two-round assessment (preliminary verification, final approval)

#### Expected contributions:

- A curated library of polyvocal ODPs tailored to the medical (e.g., specialization-based) domain, enabling reusable modules for stakeholder-specific interpretations.
- An LLM-based pipeline for ontology module creation, supporting both novel polyvocal schema generation and retrofitting of existing univocal ontologies.

These outputs provide the foundational schemas required for RQ2’s perspective-aware extraction, demonstrating polyvocality’s extensibility beyond cultural heritage to high-stakes knowledge representation.

### 3.2 RQ2: Perspective-Aware Knowledge Extraction

Given polyvocal ontologies from RQ1, this question operationalizes perspective-aware extraction

at scale. Knowledge extraction traditionally populates ontologies with structured facts from unstructured sources, yet current approaches suffer from schema inconsistency and single-perspective bias. LLMs advance this through ontology-guided prompting, but no existing frameworks systematically extract perspective-aware facts—where conditioning on polyvocal schemas yields alternative, evidence-grounded facts from identical sources while preserving epistemic separability.

RQ2 decomposes into three sub-questions addressing distinct technical challenges:

- **RQ2.1 (Agent Role Assignment):** How should perspective-conditioned LLM-based MAS be guided by polyvocal ontologies such that agents extract coherent perspective-specific triples from the same source text?
- **RQ2.2 (Provenance and Epistemic Separability):** How should provenance encode “who said what, under which perspective” to enable queryability, auditability, and conflict-aware reasoning across competing stakeholder commitments?
- **RQ2.3 (Quantitative Advantages):** What are the measurable improvements in extraction precision, perspective-specific recall, and schema adherence when using ontology-grounded MAS versus baseline single-agent approaches?

#### Technical Architecture:

- **Agent Role Assignment:** Each agent receives a perspective identity (e.g., clinical vs. epidemiological) derived from RQ1 ontologies, operationalized via perspective-specific competency questions and predicate constraints. Agents are prompted to extract only facts relevant to their assigned stakeholder viewpoint.
- **Parallel Extraction and Epistemic Separability:** Agents operate independently on identical texts, producing perspective-separated triple stores. No voting, merging, or consensus mechanisms collapse perspectives; instead, conflicting assertions remain explicit and queryable.
- **Provenance Encoding:** Each extracted triple is tagged with comprehensive provenance, enabling queries such as “Show me the interpretations of this symptom-diagnosis relationship given a certain medical specialization.”

Figure 1 illustrates the target technical architecture for perspective-aware knowledge extraction. The diagram can represent the following case: a series of medical texts (e.g., patient information from clinical visits to different medical specialties) is routed to multiple perspective-specific agents, each conditioned by a polyvocal ontology that operationalize stakeholder-specific reasoning. For instance, given a text describing a patient’s symptoms and prescribed treatment, a *clinical agent* extracts treatment-focused relationships (e.g., “azithromycin treats pneumonia”), while an *epidemiological agent* extracts population-level disease patterns (e.g., “respiratory illness increases in winter months”).

Each extracted triple is stored in perspective-specific triple stores within the polyvocal knowledge graph, tagged with comprehensive provenance metadata (perspective identity, source span, agent identifier, confidence score), enabling auditable reasoning. This design supports three classes of queries: (1) perspective-specific queries (“What clinical treatments apply?”), (2) cross-perspective comparisons (“Where do clinical and epidemiological interpretations converge or diverge?”), and (3) conflict detection (“Which facts are disputed across perspectives?”). The result is a queryable polyvocal knowledge graph that preserves competing stakeholder commitments while maintaining coherence and audibility, i.e., the core objective of this research.

**Evaluation:** On private medical datasets (and potentially public datasets), we measure:

- **Extraction Precision:** Accuracy of extracted triples against domain expert annotations for perspective-specific facts
- **Perspective-Specific Recall:** Coverage of stakeholder-relevant facts
- **Schema Adherence:** Compliance with perspective-specific predicate constraints and logical consistency within each perspective

It is important, in the evaluation phase of RQ2, to compare the proposed solution against modern Long-Context LLM or RAG baselines, since one of the justifications for the overhead of the MAS/Ontology system is the possible quantitative advantage in these tasks.

**Expected contributions:**

- A perspective-conditioned extraction pipeline leveraging polyvocal ontologies for schema-consistent, multi-perspective knowledge graph population from shared sources.
- An ontology-orchestrated multi-agent framework with explicit provenance tracking for auditable, queryable polyvocality in medical diagnosis.

These outputs establish systematic perspective-aware knowledge graph construction, extending beyond cultural heritage to high-stakes domains where competing yet valid interpretations must remain distinct, coherent, and actionable.

### 3.3 RQ3: Semantic Diversity and Schema Coherence

RQ3 is designed as a validation whether the multi-agent extraction system developed in RQ2 delivers on its promise: that extracted perspectives are genuinely diverse (not superficially relabeled duplicates), that schema constraints are maintained (not violated in pursuit of diversity), and that the MAS operates coherently across stakeholder viewpoints.

LLMs, as a “superposition of perspectives,” can generate diverse perspectivized texts when properly prompted. Hayati et al. (2024) demonstrate criteria-based prompting for extracting stances, rationales, and value-driven criteria (e.g., autonomy, teamwork), though without grounding in knowledge engineering principles. This work explores whether extraction outputs conditioned on explicit perspectivization knowledge structures (from RQ1 ontologies) enhance both diversity and schema coherence.

RQ3 decomposes into two sub-questions:

- **RQ3.1 (Semantic Diversity):** Are perspective-specific extractions from RQ2 semantically distinct, or do they represent superficial label variation of identical facts?
- **RQ3.2 (Schema Coherence):** Do perspective-conditioned extractions remain compliant with the constraints set up by the polyvocal ontological structures?

**Evaluation:** We validate RQ2 outputs against two complementary datasets. On medical data (collected during RQ2 evaluation), we measure grounding by verifying that each perspective-specific extraction is explicitly supported by the source text

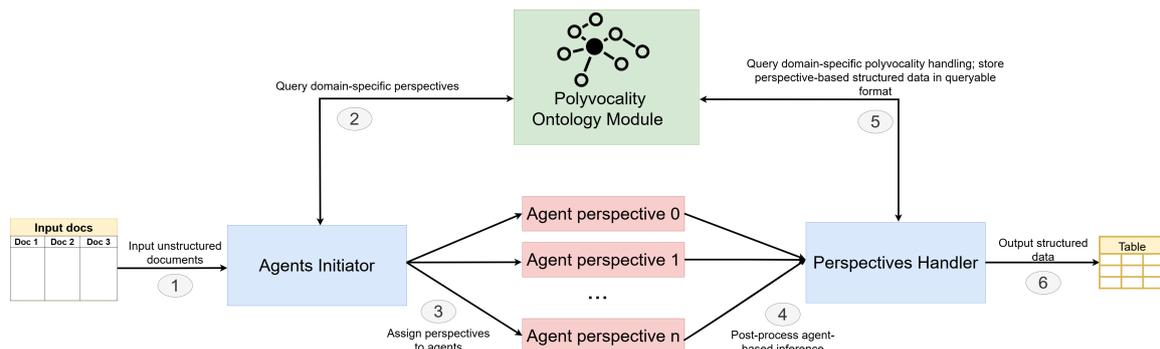


Figure 1: End-to-End Perspective-Aware Knowledge Extraction Pipeline. INPUT: A series of texts to be routed to several perspective-specific agents, each conditioned by polyvocal ontologies. Agents extract in parallel with no merging or voting, maintaining epistemic separability. Each extracted triple is stored in the polyvocal ontology enabling queryable, auditable, and conflict-aware reasoning across stakeholder viewpoints. OUTPUT: Final structured data is returned.

using Natural Language Inference (NLI). We additionally evaluate semantic diversity through the average pairwise cosine distance between perspective-specific extractions, compared against single-agent baseline extractions. Finally, we assess schema coherence by verifying that perspective-specific predicates are drawn from the appropriate ontological sub-schema and that no logical contradictions arise within a single perspective.

#### Expected contributions:

- Empirical evidence that RQ2’s framework produces genuinely diverse, non-redundant perspectives rather than copies of identical facts.
- Validation that multi-agent extraction preserves knowledge graph integrity: ontological constraints are maintained, logical consistency is ensured within perspectives, and inter-perspective coherence is auditable via provenance.

These outputs demonstrate that perspective-aware knowledge extraction is both feasible and principled—yielding queryable, stakeholder-relevant interpretations while upholding the structural and semantic rigor required for high-stakes knowledge representation in the medical domain.

## 4 Anticipated Challenges

While this research agenda outlines concrete pipelines and evaluation strategies, several anticipated challenges may constrain feasibility or require methodological adaptation. Furthermore, it is worth emphasizing that such a system must be built iteratively, adapting the schemas and patterns while testing for their usefulness in the pipeline.

One of the main challenges of this research is the generation of coherent ontology modules. LLMs may fail to generate ontology design patterns that are internally coherent or meaningfully distinguish perspectives. Generated ODPs might conflate clinical and epidemiological predicates semantically, violate ontological principles (circular dependencies, inconsistent hierarchies), or introduce contradictions when retrofitting existing schemas. If LLM-driven generation yields poor-quality ontologies, we adopt a hybrid human-in-the-loop strategy: domain experts establish gold-standard perspective distinctions through workshops, and LLMs refine rather than generate ODPs. Formal ontology validation pipelines (OWL reasoners, SPARQL competency questions) enforce coherence. This maintains RQ2’s feasibility with increased manual effort.

Another challenge is the sharing of data for results reproduction. While sharing the private datasets will not be possible, we mitigate this challenge through three strategies: releasing open-source code and evaluation scripts enabling community application to private datasets; attempting to use public datasets where applicable; and finally creating a public synthetic benchmark dataset from case studies and published guidelines, if necessary.

Finally, an orthogonal research question concerns evaluating the extent to which a given LLM can give rise to an agent that successfully adheres to a pre-defined perspective. While this question lies outside the main focus of this thesis, preliminary work will be undertaken to obtain initial evidence regarding such a kind of evaluation.

By anticipating such challenges, we aim to make RQ1, RQ2, and RQ3 remain feasible even in chal-

lenging scenarios. Should fundamental obstacles emerge, the thesis contributes by identifying these barriers and proposing principled mitigation strategies for future work.

## 5 Conclusion

This thesis proposes a systematic investigation of how polyvocal ontologies and LLM-based MAS can jointly operationalize perspective-aware knowledge extraction, thus moving beyond today’s consensus-oriented, single-perspective knowledge infrastructures toward epistemically rich, queryable representations that preserve competing stakeholder interpretations.

The research unfolds through three interconnected, sequentially dependent contributions. First (RQ1), we establish that polyvocal ontology design patterns can be systematically generated and adapted for high-stakes domains beyond cultural heritage, providing the foundational schemas necessary for perspective-conditioned extraction. Second (RQ2) we develop an ontology-grounded multi-agent framework that extracts perspective-specific facts while maintaining schema coherence, provenance traceability, and epistemic separability. Unlike existing approaches that merge or aggregate conflicting perspectives, our framework preserves them as distinct, queryable assertions traceable to stakeholder viewpoints. Third (RQ3), we validate that the resulting extractions are genuinely diverse (not superficial label variation) and maintain both schema integrity and logical consistency within perspectives.

By grounding perspective diversity in explicit ontological structures and operationalizing multi-agent coordination through provenance tracking, this work addresses a critical gap: the lack of principled mechanisms for perspective-aware knowledge representation in domains where competing interpretations matter. Evaluation on medical data—in collaboration with domain experts—demonstrates extensibility beyond cultural heritage to high-stakes applications where clinical and epidemiological viewpoints must coexist without flattening into consensus.

The anticipated outcomes are threefold: (1) a curated library of polyvocal ontology design patterns for the medical; (2) a complete framework for perspective-aware knowledge extraction with auditable provenance; and (3) empirical validation that perspective diversity is both achievable and

compatible with knowledge graph rigor. Together, these contributions establish the technical and conceptual foundation for building fairer, more transparent AI systems that honor the epistemic diversity inherent in complex domains while maintaining interoperability and auditability.

This research thus demonstrates that polyvocality can become a first-class principle for knowledge representation in healthcare, science, policy, and beyond.

## Limitations

This research agenda, while outlining concrete pipelines and benchmarks, remains at the proposal stage without full empirical implementation across all sub-questions. Evaluations rely on private datasets whose availability may limit reproducibility, though ethical handling and expert collaboration mitigate privacy risks.

The focus on the medical domain, while extending polyvocality beyond cultural heritage, does not yet address other high-stakes fields like policy or manufacturing, where stakeholder conflicts may require additional ODP adaptations.

## Acknowledgments

This publication was funded by a flagship project “CHEXRISH: Cultural Heritage Exploration and Retrieval with Intelligent Systems at Jagiellonian University” under the Strategic Programme Excellence Initiative at Jagiellonian University.

The research for this publication has been supported by a grant from the Priority Research Area DigiWorld under the Strategic Programme Excellence Initiative at Jagiellonian University.

During the preparation of this work, the authors used GPT-5.1 and GPT-4o in order to: grammar and spelling check, paraphrase and reword. After using these services, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

## Ethics statement

The project involves sensitive patient data. Therefore, to conduct the experiments, necessary ethics approval will be sought from the Ethics Committee of the host university. We will limit the research to models that can fit in local machines to avoid external API data transmission risks, ensuring strict anonymization.

## References

- Federico Cabitza, Andrea Campagner, and Valerio Basile. 2023. [Toward a perspectivist turn in ground truthing for predictive computing](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6):6860–6868.
- Simona Frenda, Gavin Abercrombie, Valerio Basile, Alessandro Pedrani, Raffaella Panizzon, Alessandra Teresa Cignarella, Cristina Marco, and Davide Bernardi. 2024. [Perspectivist approaches to natural language processing: a survey: Perspectivist approaches to natural language processing...](#) *Lang. Resour. Eval.*, 59(2):1719–1746.
- Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. [Bias and fairness in large language models: A survey](#). *Computational Linguistics*, 50(3):1097–1179.
- Aldo Gangemi and Valentina Presutti. 2022. Formal representation and extraction of perspectives. In Piek Vossen and Antske Fokkens, editors, *Creating a More Transparent Internet*, pages 208–228. Cambridge University Press, Cambridge.
- Alireza Ghafarollahi and Markus J. Buehler. 2025. [Sciagents: Automating scientific discovery through bioinspired multi-agent intelligent graph reasoning](#). *Advanced Materials*, 37(22):2413523.
- Shirley Anugrah Hayati, Minhwa Lee, Dheeraj Rajagopal, and Dongyeop Kang. 2024. [How far can we extract diverse perspectives from large language models?](#) In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5336–5366, Miami, Florida, USA. Association for Computational Linguistics.
- Alistair Johnson, Tom Pollard, and Roger Mark. 2016. [MIMIC-III clinical database \(version 1.4\)](#). PhysioNet.
- Jackie Kay, Atoosa Kasirzadeh, and Shakir Mohamed. 2025. [Epistemic Injustice in Generative AI](#), page 684–697. AAAI Press.
- Samira Khorshidi, Azadeh Nikfarjam, Suprita Shankar, Yisi Sang, Yash Govind, Hyun Jang, Ali Kasgari, Alexis McClimans, Mohamed Soliman, Vishnu Vardhan Reddy Konda, Ahmed Fakhry, and Xiaoguang Qi. 2025. [Odke+: Ontology-guided open-domain knowledge extraction with llms](#). *ArXiv*, abs/2509.04696.
- Grgur Kovač, Masataka Sawayama, Rémy Portelas, Cédric Colas, Peter Ford Dominey, and Pierre-Yves Oudeyer. 2023. [Large language models as superpositions of cultural perspectives](#). *ArXiv*, abs/2307.07870.
- Angelie Kraft and Eloïse Soulier. 2024. [Knowledge-enhanced language models are not bias-proof: Situated knowledge and epistemic injustice in ai](#). In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24*, page 1433–1445, New York, NY, USA. Association for Computing Machinery.
- Anna Sofia Lippolis, Mohammad Javad Saeedizade, Robin Keskiärrkkä, Sara Zuppiroli, Miguel Ceriani, Aldo Gangemi, Eva Blomqvist, and Andrea Giovanni Nuzzolese. 2025. [Ontology generation using large language models](#). In *The Semantic Web: 22nd European Semantic Web Conference, ESWC 2025, Portoroz, Slovenia, June 1–5, 2025, Proceedings, Part I*, page 321–341, Berlin, Heidelberg. Springer-Verlag.
- Sanaz Saki Norouzi, Adrita Barua, Antrea Christou, Nikita Gautam, Andrew Eells, Pascal Hitzler, and Cogan Shimizu. 2024. [Ontology population using llms](#). *Preprint*, arXiv:2411.01612.
- Sanghyun Park, Boris Maciejovsky, and Phanish Puranam. 2025. [Thinking with many minds: Using large language models for multi-perspective problem-solving](#). *Preprint*, arXiv:2501.02348.
- Sami Saadaoui and Eduardo Alonso. 2025. [Coordinated llm multi-agent systems for collaborative question-answer generation](#). *Knowledge-Based Systems*, 330:114627.
- Cogan Shimizu and Pascal Hitzler. 2025. [Accelerating knowledge graph and ontology engineering with large language models](#). *Journal of Web Semantics*, 85:100862.
- Sarah Binta Alam Shoilee, Victor de Boer, and Jacco van Ossenbruggen. 2023. [Polyvocal Knowledge Modelling for Ethnographic Heritage Object Provenance](#), pages 127–143. Studies on semantic web. IOS Press.
- Marieke van Erp and Victor de Boer. 2021. A polyvocal and contextualised semantic web. In *The Semantic Web*, pages 506–512, Cham. Springer International Publishing.
- Bohui Zhang, Valentina Anita Carriero, Katrin Schreiberhuber, Stefani Tsaneva, Lucía Sánchez González, Jongmo Kim, and Jacopo de Berardinis. 2025. [Ontochat: A framework for conversational ontology engineering using language models](#). In *The Semantic Web: ESWC 2024 Satellite Events*, pages 102–121, Cham. Springer Nature Switzerland.
- Kaiwen Zuo, Yirui Jiang, Fan Mo, and Pietro Lio. 2024. [Kg4diagnosis: A hierarchical multi-agent llm framework with knowledge graph enhancement for medical diagnosis](#). *ArXiv*, abs/2412.16833.