Multi-Perspective Ontology Alignment: Bridging Epistemic Differences in a Water Knowledge Case Study

Divyasha Sunil Naik[1,*], Birgitta König-Ries[1]

[1]Friedrich Schiller University Jena, Germany

Abstract

Ontologies offer structured representations of domain knowledge, but their conceptualizations often reflect the goals of specific disciplinary standpoints. In complex domains like water, where meanings span from chemical substance to cultural symbol, aligning these divergent ontologies presents a significant challenge. Existing methods often overlook deeper conceptual mismatches. This paper introduces a methodology for multi-perspective ontology alignment that preserves epistemic diversity while enabling structured exploration. The approach models conceptual overlaps and tensions across disciplines through standpoint tagging, bridge relation discovery, and formal representation. We further explore how AI agents, including LLM-based systems, can support analogical reasoning and suggest new conceptual links. Using water-related ontologies as a case study, we demonstrate how our framework enables cross-perspective querying and semantic navigation without forcing ontological convergence.

Keywords

Semantic interoperability, Knowledge representation, Epistemic perspectives, Large language models, Water ontologies

1. Introduction

Natural language terms rarely possess precise or universally agreed-upon definitions. Their meanings shift depending on context, situation, and the perspective of the speaker. Ambiguity often arises when and to what extent a term applies in a given scenario. In such borderline cases, speakers and, by extension, ontology designers must make deliberate choices about how to apply each term.

When a domain is formalized into an ontology, the ontology explicitly encodes these choices as conceptual commitments that shape its structure and scope. Each ontology thus becomes a reflection not only of the modeled domain, but also of the worldview, assumptions, and design objectives of those who construct it. As a result, differences in intended granularity, scope, and perspective naturally emerge, meaning that distinct communities may construct ontologies that diverge significantly—even when representing the same real-world phenomenon.

Without the ability to bridge across perspectives, interdisciplinary problems remain fragmented, limiting both understanding and practical intervention. This motivates the need for structured methods that preserve conceptual diversity while enabling semantic connection and exploration. However, perspectival divergence presents a critical challenge for aligning ontologies across disciplinary boundaries. Multiple valid conceptualizations of the same real-world entity, shaped by distinct epistemic standpoints, must be related to preserve their unique insights while supporting semantic interoperability. Yet most existing alignment techniques typically focus more on lexical similarity [1, 2], and often fail to capture deeper conceptual differences embedded in the ontology. Bridging these differences requires methods that represent, relate, and reason over multiple perspectives.

36th GI-Workshop on Foundations of Databases (Grundlagen von Datenbanken), September 29 - October 01, 2025, Regensburg, Germany

*Corresponding author.

divyasha.sunil.naik@uni-jena.de (D. S. Naik); birgitta.koenig-ries@uni-jena.de (B. König-Ries)

https://www.fmi.uni-jena.de/en/10912/divyasha-sunil-naik (D. S. Naik); https://www.fmi.uni-jena.de/en/10780/koenig-ries (B. König-Ries)

0009-0003-9374-6951 (D. S. Naik); 0000-0002-2382-9722 (B. König-Ries)

1.1. Problem Statement

Our collaboration within the Thuringian Water Innovation Cluster (ThWIC) [3] has highlighted the necessity of integrating diverse disciplinary insights to address the complex, interrelated challenges associated with water. The meaning and relevance of water vary dramatically across domains: as a chemical substance (H_2O) in chemistry, a physical component of rivers and ecosystems in hydrology, a regulated resource in policy, or a cultural and symbolic entity in religion and philosophy. While it is possible to construct ontologies that faithfully capture a single disciplinary perspective, such representations often fall short of modeling the full semantic breadth and practical interdependencies that real-world water issues entail. Bringing together existing water-related ontologies across these domains is particularly challenging, not only due to conceptual differences, but also because it demands significant human effort and deep domain expertise. Despite the availability of advanced tools, ontology engineering remains time-consuming and error-prone, requiring careful attention to define entities, relationships, and semantic boundaries.

These challenges highlight a critical gap: while individual ontologies may successfully represent water within a given disciplinary framework, there is no general method for systematically identifying, relating, and querying diverse conceptualizations in water-related ontologies. To address this, we explore how semantically distinct representations such as H_2O in molecular chemistry, *"water of life"* in symbolic traditions, or a *"regulated resource"* in policy can be brought into relation without forcing ontological convergence. This requires analyzing the types of relationships that exist between these conceptualizations, identifying where generalization or disambiguation is needed, and making semantic differences explicit rather than suppressing them. In doing so, we treat misalignments, conflicts, and pragmatic differences not as problems to be eliminated or resolved, but as essential features to be modeled and understood.

No comprehensive methodology yet captures perspectival diversity while maintaining each perspective's structure, meaning, and intent. This challenge is especially acute for complex domains like water, where overlapping yet non-equivalent viewpoints must coexist meaningfully. To support this, we propose a conceptual methodology for the structured exploration of heterogeneous ontologies through a multi-perspective lens—preserving the boundaries of each domain while facilitating connections where appropriate. This paper presents the intended study design and a step-by-step methodology for aligning heterogeneous ontologies from multiple epistemic perspectives. Figure 3 and Figure 4 illustrate the envisioned structure and bridging relationships that this method aims to support across heterogeneous ontologies.

2. Related Work

Ontology alignment addresses semantic heterogeneity by identifying correspondences between entities across knowledge sources. However, most methods focus on semantic or lexical similarity, with limited support for epistemic or perspectival differences.

Traditional tools such as LogMap [4], COMA [5], and AgreementMaker [6] focus on identifying correspondences based on lexical or logic-based similarity, typically assuming a single, unified conceptualization of the domain. Machine learning methods like BERTMap [7] and LogMap-ML [8] apply NLP-based techniques to improve scalability and reduce manual effort.

LLM-based approaches have introduced prompt-based classification [9, 10] and embedding-based alignment, as demonstrated by methods such as OLaLa [11] and LLMs4OM [12], to support entity matching across ontologies.

MILA [13] achieves high accuracy and efficiency in ontology matching with a retrieve-identify-prompt pipeline, but its focus remains primarily on pairwise correspondences and does not fully address more complex, higher-order relationships. Similarly, OntoAligner [14] focuses on entity-level and semantic similarity-based alignments, yet does not explicitly address complex structural correspondences. Sousa et al. [15] address complex mappings using LLM embeddings, although their method remains computationally intensive and narrowly scoped.

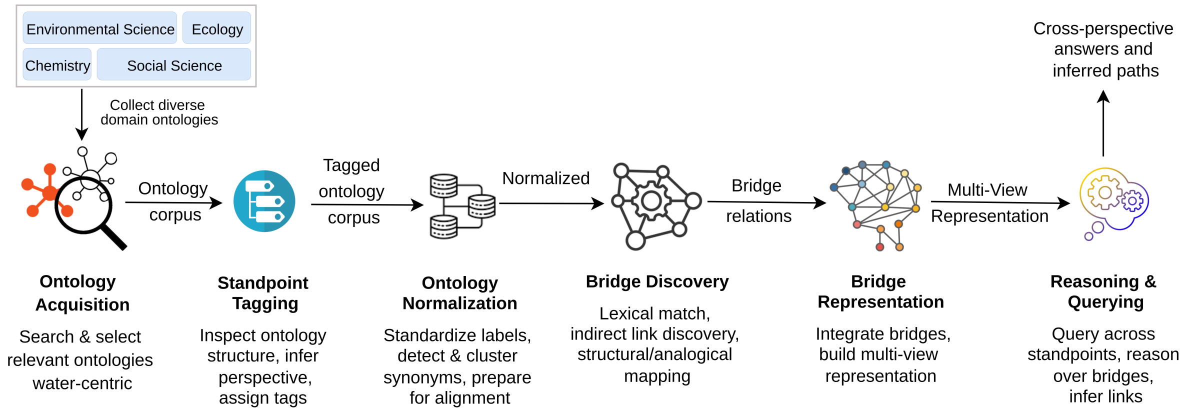

Figure 1: Workflow of the proposed multi-perspective ontology alignment approach, showing the sequential stages and intermediate artifacts described in Section 3.

Multi-viewpoint ontology alignment has also been explored by Kolli et al. [16], who propose a rule-based approach using reaction rules to compose mappings across perspectives. However, their method depends on pre-existing correspondences and does not support the discovery of bridges between previously unaligned or semantically distant concepts.

In contrast, we propose a conceptual methodology for discovering indirect and analogical bridges — drawing inspiration from structure-mapping theory [17] and indirect alignment [18] to relate structurally similar but epistemically distinct concepts. We further outline a vision for an agent-navigable representation that would support viewpoint-aware reasoning and querying, informed by frameworks such as AGNO [19]. Our methodology aims to address the challenge of multi-perspective modeling by preserving conceptual boundaries while enabling structured navigation across disciplinary views.

3. Methodological Steps

This chapter presents a step-by-step methodology to model, represent, and bridge the conceptual relationships, overlaps, and tensions that arise when different communities describe the same real-world phenomenon, such as water, through distinct epistemic frameworks, using ontologies as structured representations. The methodology is organized into seven sequential stages that progressively build a structured, multi-view representation connecting diverse conceptualizations, as described in Figure 1. The following subsections discuss each stage in detail.

- **Stage 1:** Entails the collection of ontologies from diverse domains, focusing on a shared theme (e.g., water).
- **Stage 2:** Introduces a standpoint tagging process to label each ontology or concept according to its disciplinary perspective (e.g., scientific, cultural, policy-related).
- **Stage 3:** Ensures cross-ontology normalization, where concept labels are standardized and similar terms clustered to enable later alignment.
- **Stage 4:** Focuses on discovering bridge relations between concept nodes across perspectives, through analogical and indirect techniques.
- **Stage 5:** Covers the construction of a bridge ontology or multi-view representation that formally links these concepts, illustrated in Figure 3 and Figure 4
- **Stage 6:** Applies reasoning and querying over this structure to demonstrate its usefulness.
- **Stage 7:** Introduces an AI agent capable of navigating, extending, or explaining these conceptual bridges, supporting partial automation and intelligent exploration.

3.1. Stage 1: Ontology Acquisition

The first step to establishing a robust foundation involves systematically collecting a diverse set of water-related ontologies from multiple disciplines, such as ecology, chemistry, environmental science, and the social sciences. This ensures broad coverage of conceptualizations that reflect the multifaceted nature of water.

Established ontology repositories, including NCBI BioPortal, are explored to identify ontologies or relevant modules within larger ontologies that conceptualize water. The focus is placed on concepts and entities that are either explicitly or implicitly related to "water." This includes direct instances (e.g., H_2O, river, drinking water) as well as semantically related subclasses, roles, and contextual definitions that reflect diverse domain-specific perspectives. The ontology corpus remains open to iterative extension as additional relevant sources are discovered during the course of the study.

3.2. Stage 2: Standpoint Tagging of Ontologies

Once the ontology corpus has been assembled, the next step will involve assigning each ontology or its relevant water-related modules a **standpoint tag** that will reflect the underlying disciplinary or epistemic perspective from which the concept of water is modeled. These tags will help to preserve the diversity of viewpoints across domains and will enable structured comparison without suppressing viewpoint-specific differences.

Each ontology will be tagged based on a combination of indicators, including its top-level class hierarchy, the definitions or comments associated with key water-related entities, and the domain and range of relevant properties. For example, an ontology that defines water as a subclass of Molecule and includes terms such as H_2O, *chemical compound*, or *solvent* will likely be assigned a **Scientific** tag. In contrast, ontologies referring to water in the context of *policy instruments*, *governance frameworks*, or *sustainable development goals* will be tagged as **Policy-Oriented**. Similarly, ontologies referencing water in symbolic, religious, or ritual contexts will be marked as **Cultural**, while those emphasizing learning, training, or community engagement may be labeled as **Educational**.

This tagging process will initially be performed manually for a curated set of ontologies to ensure semantic precision. However, a semi-automatic procedure may also be used to support scaling. This procedure will leverage a predefined dictionary of trigger terms and simple NLP techniques to suggest likely tags based on class names, comments, and property labels. The standpoint tags will serve as an essential input to later stages of alignment, enabling targeted identification of conceptual overlaps, conflicts, and possible bridges between distinct perspectives on water.

3.3. Stage 3: Cross-Ontology Normalization

To prepare for meaningful alignment, the collected and tagged ontologies will undergo a normalization step aimed at minimizing surface-level differences and identifying semantically equivalent or related concepts across perspectives. This stage will ensure that terminological inconsistencies, such as varying naming conventions or minor syntactic differences, do not hinder the discovery of actual conceptual relationships.

1. **Label Standardization:**
   Class and property names across the ontologies will be cleaned and standardized to follow a uniform format. This will include converting labels to lowercase, removing underscores or camel case, expanding abbreviations, and ensuring consistent use of spacing or punctuation. The goal will not be to alter the ontology's semantics, but to enable reliable matching and comparison.
2. **Synonym Detection and Clustering:**
   After standardization, concept labels will be analyzed to detect synonyms and near-synonyms across ontologies. This step will use lexical similarity measures such as cosine similarity over TF-IDF vectors and word embeddings (e.g., Word2Vec). Concept labels that refer to similar or overlapping ideas will be grouped accordingly. For instance, concept nodes such as *HolyWater*

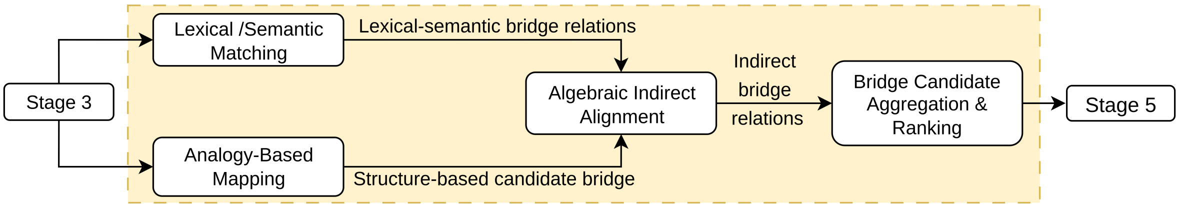

Figure 2: Workflow for conceptual bridge discovery across multiple perspectives (Stage 4)

and *BlessedWater* may be grouped together based on their ritual and symbolic significance. Similarly, *H_2O*, *WaterMolecule*, and *LiquidWater* can be clustered due to their scientific and physical definitions. These groupings will help reveal implicit similarities despite differing terminologies.

3.4. Stage 4: Discovery of Bridge Relations Across Perspectives

This stage will operationalize the core objective of bridging epistemic divergences in multi-perspective ontology alignment by discovering conceptual bridges. These bridges will form the semantic backbone connecting siloed ontologies while preserving their distinct epistemic perspectives. The input to this stage will consist of normalized ontologies enriched with epistemic standpoint tags and standardized concept labels and relationships, as prepared in Stage 3. Figure 2 will illustrate the workflow of this stage, showing how lexical-semantic matching and analogy-based mapping will feed into algebraic indirect alignment, which will produce indirect bridge relations for candidate aggregation and ranking before proceeding to Stage 5.

- **Lexical Matching:** Concepts with similar surface labels—such as *WaterGovernance* and *WaterRegulation*—will be aligned using string similarity techniques and synonym expansion through resources like WordNet. This lexical similarity process will detect candidates based on both direct string matching and semantic closeness, including synonyms and hypernyms.
- **Structural / Analogical-Based Mapping:** Operating in parallel with lexical matching, this step will employ a cognitively inspired analogy-based alignment approach modeled after the Structure-Mapping Theory (SMT) [17] framework. Each ontology will be represented as a relational graph in which concepts and their interconnections form patterns of semantic roles. Using a process akin to human analogical reasoning, the method will detect correspondences between concepts that occupy analogous roles. Activation will spread between concept nodes based on shared semantic features and relational context, enabling the identification of structural bridges that reveal functional or contextual equivalences. For instance, while *HolyWater* in a religious ontology may be involved in `usedIn:Ritual`, and *WaterLiteracy* in an educational ontology may relate to `usedIn:AwarenessCampaign`, both will be interpreted as culturally meaningful interventions.
- **Algebraic Indirect Alignment:** Following the parallel lexical matching and analogy-based mapping steps, this stage will address the challenge of bridging conceptual gaps where direct alignments are incomplete or absent. It will employ an algebraic composition approach, grounded in alignment algebra, to infer indirect alignments by chaining existing direct mappings through intermediate ontologies or concepts [18]. When concepts do not match directly, connections will be inferred via shared relations to an intermediate node.
  For instance, *HolyWater* (Cultural) and *WaterAwarenessProgram* (Educational) may both relate to a concept like *RitualPractice* or *PublicEngagement*, thereby enabling indirect alignment. This method will consider both the semantics of relationships and the confidence values associated with each direct alignment. By reusing and composing existing alignments through mathematical rules, it will generate additional candidate mappings that will extend coverage and support the connection of epistemically diverse perspectives within the multi-perspective alignment framework.

- **Bridge Candidate Aggregation and Ranking:** This step will aggregate candidate bridges identified in the previous stages—lexical, analogical, and indirect. It will implement a ranking scheme to prioritize the most robust and epistemically justified alignments. The ranking will be guided by combined confidence scores along with measures of semantic coherence and structural consistency, ultimately producing a prioritized list of reliable cross-ontology bridges for effective multi-perspective integration.

To support transparency and future automation, each discovered bridge will be annotated with metadata, such as the type of relation (e.g., `is_contextualized_by`, `engages_with`, `structurally_analogous_to`) and its provenance (e.g., lexical match, indirect chain, analogical inference). The output of this stage will be a documented set of cross-perspective alignments that will preserve contextual meaning while enabling structured linkage.

3.5. Stage 5: Formal Representation of Multi-Perspective Connections

The discovered bridge relations will be formalized into a structured representation that will link concepts across disciplinary standpoints. This will result in a bridge ontology that encodes the original concepts and the semantic paths connecting them. The construction process proceeds by integrating:

- **Original standpoint-specific nodes:** Concepts from the source ontologies will be preserved as distinct nodes, retaining their domain-specific definitions and tags (e.g., *HolyWater* Cultural, *H_2O* Scientific, *Water Used in Rituals* Social).
- **Bridge nodes and edges:** Intermediate concepts (e.g., *RitualPractice*, *ContextualizedWaterUse*) and discovered bridge relations (e.g., *engages_with*, *is_contextualized_by*) will be introduced as connectors between existing nodes.
- **Semantic qualifiers:** Where needed, edges will be annotated with qualifiers to indicate the source of the bridge (e.g., analogical, indirect, structural) or its level of certainty.

Figure 4 illustrates an example of such a multi-perspective structure. Concepts like *HolyWater* and *Water in Discourse*, though originating from different ontologies and perspectives, will now be connected via intermediate notions such as *has_cultural_context* or *shaped_in*, reflecting shared functions or social roles. To improve semantic clarity and interoperability, bridge relations identified in the previous stages will be mapped to existing vocabularies and ontologies whenever possible. Instead of introducing custom terms, commonly used relation ontologies such as RO (Relation Ontology), SKOS, OWL object properties, or upper ontologies like BFO and DOLCE are reused. Alignment properties such as `owl:sameAs`, `skos:exactMatch`, or `schema:about` are applied where appropriate. New properties will be introduced only when no suitable relation exists. This step ensures that the resulting representation can support reasoning, link with external datasets, and remain compatible with semantic web standards. The resulting structure will preserve the distinct meanings and contexts of each original concept, while enabling navigation and reasoning across perspectives. It will form the basis for downstream tasks such as multi-standpoint querying, cross-domain explanation, and automated exploration by AI agents.

3.6. Stage 6: Cross-Perspective Reasoning and Querying

With the multi-perspective representation established in Stage 4, this stage will demonstrate how the resulting structure can be queried and reasoned over to support cross-disciplinary exploration. The goal will be to validate the usefulness of the bridge ontology as both a conceptual and queryable structure. Queries will be designed to navigate across standpoints, retrieve interconnected concepts, and reveal hidden semantic paths between perspectives. In addition to basic querying, the structure supports simple reasoning tasks. For instance, given a chain of bridge relations across three standpoints (e.g., cultural ⟶ ritual ⟶ educational), a reasoner will be able to infer transitive or contextual relevance of distant nodes. Semantic qualifiers guide reasoning depth and scope, enabling filtered inference across perspectives. This stage will demonstrate that the alignment is not merely visual or conceptual—it will

Table 1
Definitions of concept nodes in Figure 1, provided because they are not found in existing ontologies.

Node label	Definition
Water	A fundamental natural substance essential for life, found in various contexts, e.g., physical, social, symbolic, etc
Holy Water	Water that has been blessed and is used in spiritual or religious rituals as a sacred substance.
Ritual Practice	A culturally or spiritually significant action or series of actions performed in a set sequence, often symbolizing belief or tradition.
Drinking Water	Water that is safe and intended for human consumption.
Polluted Water	Water that has been contaminated by harmful substances and is no longer safe for its intended use.
Pollutant	Any chemical, physical, or biological agent that causes pollution in air, water, or soil.
Policy Linked Water	Conceptualization of water as it is governed, regulated, or influenced by public policy, legal frameworks, or institutional guidelines.
Water Resources	Naturally occurring or managed supplies of water used for human, agricultural, ecological, or industrial purposes.

enable structured access to distributed meanings, support human interpretation, and set the stage for intelligent navigation by agents.

3.7. Stage 7: AI Agent Integration and Augmented Exploration

In the final stage, an AI agent will be integrated to explore the multi-perspective ontology. Drawing on earlier components—standpoint-tagged ontologies, clusters, and bridge relations—the agent will perform standpoint-aware querying, propose new links, and explain existing alignments. Techniques include SPARQL navigation and LLM-driven analogical inference. Agent-based frameworks like AGNO [19] and OntoAgent will support structured traversal and bridge discovery, advancing automation and scalability.

Figure 3 and Figure 4 are idealized diagrams constructed to illustrate the envisioned outcome of the proposed methodology. They do not directly originate from existing ontologies but demonstrate how concepts from multiple perspectives can be structured and interconnected. Table 1 and Table 2 provide definitions of the concept nodes and edge labels used in Figure 3, supporting clarity in interpreting the modeled perspectives and their connections. Figure 3 illustrates how different disciplines think about water. At the center is the concept of "Water," from which various perspectives are derived, such as water in religious rituals, water in policy, water in digital models, water as a natural resource, and so on. Each of these perspectives comes from a different domain and reflects how that field defines or uses the idea of water.

Figure 4 expands this view by showing how these disciplinary concepts are connected. It models conceptual bridges like shaped_in that demonstrate how water-related ideas influence one another across domains. For example, how polluted water leads to the use of water indicators, which reveal patterns in water consumption, which can then shape water policy. As an outcome, the system will be able to address questions such as:

- How does cultural perception of water influence ritual use?
- Trace the impact of pollution on governance.
- What digital models are influenced by policy and used in planning?

These examples illustrate the type of cross-perspective querying and conceptual navigation the proposed methodology is designed to support.

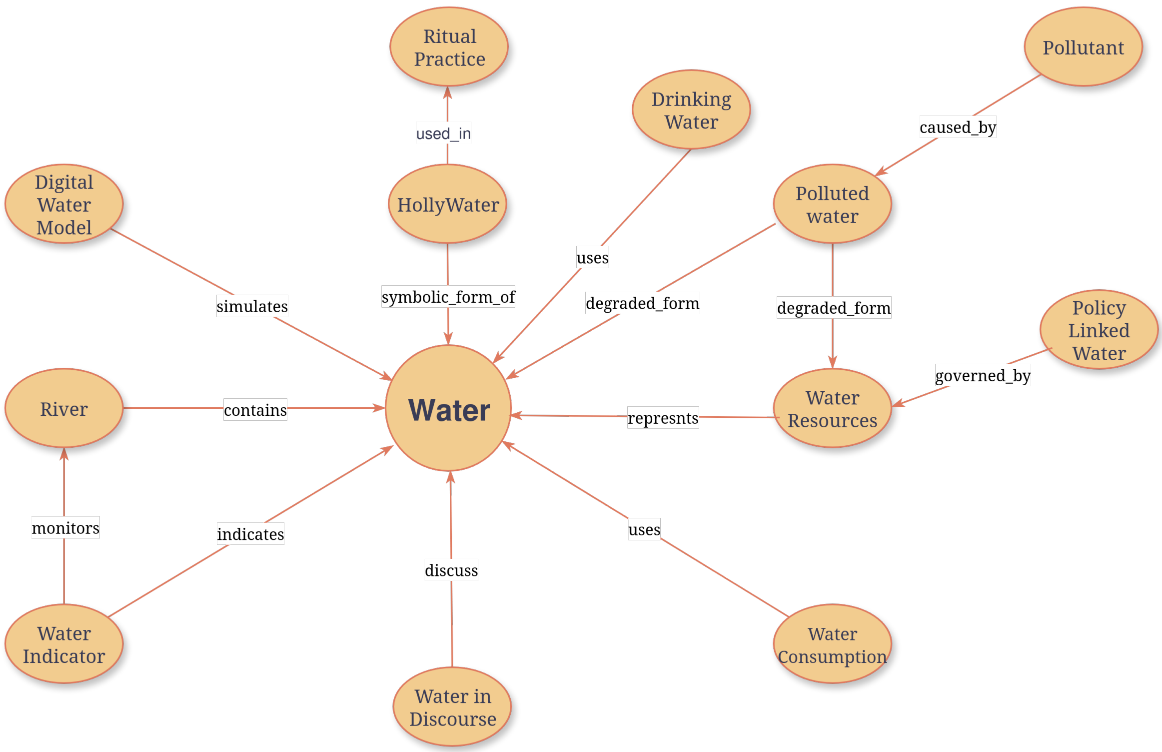

Figure 3: Shows the initial ontology modules, each from different disciplinary standpoints. Nodes like HolyWater, PolicyLinkedWater, or DigitalWaterModel represent how 'Water' is understood across domains.

Table 2
Definitions of relationship labels (edges) in Figure 1, explaining how concept nodes are connected and what each connection means.

Edge Label	Definition
used_in	Indicates that a concept or object is employed as part of an activity or context.
symbolic_form_of	Suggests that one concept symbolically represents or embodies the meaning of another.
uses	Indicates that one concept depends on or uses another for its function or purpose.
degraded_form	Indicates that one concept is a less pure or less desirable version of another due to contamination or change.
represents	One concept stands in for, reflects, or models another concept or reality.
governed by	Indicates that a concept is regulated, controlled, or influenced by another (typically policy or authority).
indicates	shows that a concept provides information or a signal about the status or presence of another.
monitors	Indicates that a concept observes or tracks the condition
contains	Indicates that one concept physically or logically includes another within its boundaries.
simulates	Shows that one concept models or structures of another in a digital
discuss	Indicates that one concept is examined, mentioned, or analyzed within the context of another.

4. Preliminary Results and Motivation

This section summarizes early experiments assessing how far existing tools address our core problem. We extracted water-centric subgraphs from a curated set of ontologies, applied LogMap as a baseline matcher, reasoned over merged subgraphs to detect logical conflicts, and used COMA++ for expert-guided alignment refinement. These exploratory efforts are not direct implementations of the methodology in

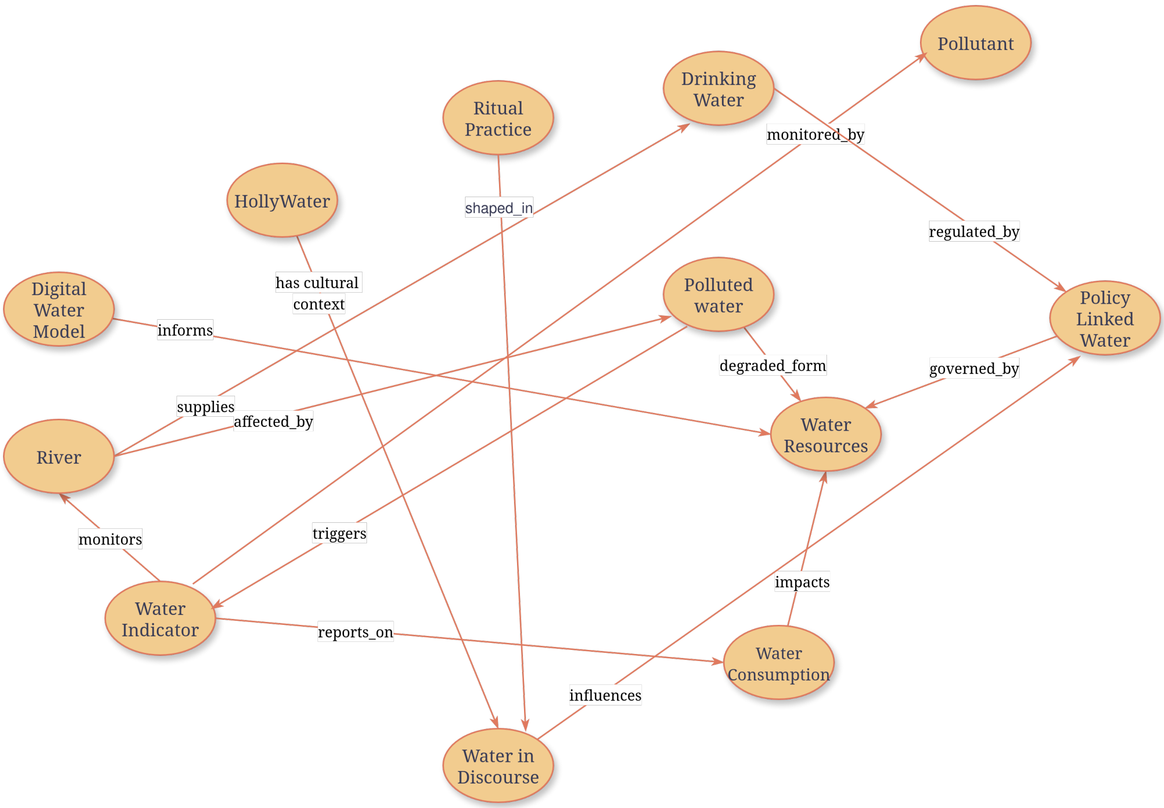

Figure 4: Illustrates the result of applying bridge discovery methodology (Step 3). New edges represent relationships between concepts across different standpoints

Section 3; rather, they are initial trials that exposed key limitations and motivated the multi-perspective approach we propose.

4.1. Ontology Survey, Selection, and Preprocessing

Initially, we surveyed existing water-related ontologies and selected five core candidates for detailed analysis: [20], [21], [22], [23], and [24]. Together, these ontologies cover complementary aspects of the water domain, including hydrological features and catchments, governance and policy frameworks, sensor-based monitoring and infrastructure, water quality observations, and data interoperability.

Because dedicated water ontologies are limited, we extended the search to BioPortal and identified 27 additional ontologies containing the term "water." The occurrence and relevance of the term varied: in some cases, water appeared as a root concept, while in others it was deeply nested. To address this variation, we extracted water-centric subgraphs—localized conceptual fragments centered on water—rather than using complete ontologies. Relevant files were downloaded in OWL format, and automated scripts retrieved ontology IDs and subtree root IDs to extract the fragments. Each resulting ontology was then loaded into Protégé for inspection and further analysis.

To ensure the corpus reflects the interdisciplinary scope of water research, we applied a curated set of water-related keywords developed from ThWIC expert input. These keywords guided ontology selection and later subgraph extraction, supporting comprehensive coverage of scientific, technical, and societal perspectives. The corpus remains open to iterative extension as additional sources are discovered.

4.2. Baseline Automated Alignment Trials

To test how well existing tools perform, we applied LogMap—a well-known ontology matcher combining lexical matching with logic-based reasoning—to the extracted fragments. LogMap processed the ontologies, generating unsatisfiable classes, but limited lexical overlap produced very few alignments, and no stable mappings or anchors were retained.

To examine structural consistency manually, we selected two semantically related subgraphs and merged them in Protégé. Reasoning with ELK and HermiT revealed conflicting subclass assertions, misaligned domain and range constraints, and overlapping or incompatible class hierarchies. As no automatic repair suggestions were generated, we refined the process using COMA++, an interactive ontology matching tool. Its configurable strategies and visual interface enabled experts to validate and adjust candidate alignments, partially mitigating inconsistencies and preparing the ontologies for integration.

These trials show that existing methods are insufficient for this domain and that expert input is often required to resolve conflicts. They motivate our current methodology, which introduces standpoint tagging, bridge discovery, formal multi-view representation, and AI-agent support to enable context-aware, scalable alignment across diverse epistemic perspectives.

5. Conclusion

This paper outlined a conceptual methodology for relating ontologies across diverse epistemic perspectives while preserving their distinctions. Using the water domain as a case, the approach combines standpoint tagging, structured representation, and bridge relation discovery to enable perspective-aware querying. AI agents can further support exploration by surfacing analogical links and guiding navigation across conceptual paths. Together, these elements demonstrate how heterogeneous viewpoints can be connected in a structured and interpretable way.

Acknowledgments

This work is supported by the German Federal Ministry of Education and Research through the ThWIC [3].

Declaration on Generative AI

During the preparation of this work, the author(s) used X-GPT-4 and Gramby in order to: Grammar and spelling check. Further, the author(s) used X-AI-IMG for figures 3 and 4 in order to: Generate images. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] M. Cheatham, P. Hitzler, String similarity metrics for ontology alignment, in: International semantic web conference, Springer, 2013, pp. 294–309.
[2] M. Yu, G. Li, D. Deng, J. Feng, String similarity search and join: a survey, Frontiers of Computer Science 10 (2016) 399–417.
[3] Thuringian Water Innovation Cluster (ThWIC), Thuringian water innovation cluster (thwic), https://www.thwic.uni-jena.de, 2025. Accessed: 2025-07-09.
[4] E. Jiménez-Ruiz, B. Cuenca Grau, Logmap: Logic-based and scalable ontology matching, in: International Semantic Web Conference, Springer, 2011, pp. 273–288.

[5] D. Aumueller, H.-H. Do, S. Massmann, E. Rahm, Schema and ontology matching with coma++, in: Proceedings of the 2005 ACM SIGMOD international conference on Management of data, 2005, pp. 906–908.

[6] D. Faria, C. Pesquita, E. Santos, M. Palmonari, I. F. Cruz, F. M. Couto, The agreementmakerlight ontology matching system, in: On the Move to Meaningful Internet Systems: OTM 2013 Conferences: Confederated International Conferences: CoopIS, DOA-Trusted Cloud, and ODBASE 2013, Graz, Austria, September 9-13, 2013. Proceedings, Springer, 2013, pp. 527–541.

[7] Y. He, J. Chen, D. Antonyrajah, I. Horrocks, Bertmap: a bert-based ontology alignment system, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, 2022, pp. 5684–5691.

[8] J. Chen, E. Jiménez-Ruiz, I. Horrocks, D. Antonyrajah, A. Hadian, J. Lee, Augmenting ontology alignment by semantic embedding and distant supervision, in: European Semantic Web Conference, Springer, 2021, pp. 392–408.

[9] Y. He, J. Chen, H. Dong, I. Horrocks, Exploring large language models for ontology alignment, arXiv preprint arXiv:2309.07172 (2023).

[10] S. Saki Norouzi, M. Saeid Mahdavinejad, P. Hitzler, Conversational ontology alignment with chatgpt, arXiv e-prints (2023) arXiv–2308.

[11] S. Hertling, H. Paulheim, Olala: Ontology matching with large language models, in: Proceedings of the 12th knowledge capture conference 2023, 2023, pp. 131–139.

[12] H. B. Giglou, J. D'Souza, F. Engel, S. Auer, Llms4om: Matching ontologies with large language models, arXiv preprint arXiv:2404.10317 (2024).

[13] M. Taboada, D. Martinez, M. Arideh, R. Mosquera, Ontology matching with large language models and prioritized depth-first search, arXiv preprint arXiv:2501.11441 (2025).

[14] H. Babaei Giglou, J. D'Souza, O. Karras, S. Auer, Ontoaligner: A comprehensive modular and robust python toolkit for ontology alignment, in: European Semantic Web Conference, Springer, 2025, pp. 174–191.

[15] G. Sousa, R. Lima, C. Trojahn, Complex ontology matching with large language model embeddings, arXiv preprint arXiv:2502.13619 (2025).

[16] M. KOLLI, Formal modeling of multi-viewpoint ontology alignment by mappings composition, Acta Univ. Sapientiae Informatica 15 (2023) 181–204.

[17] B. Falkenhainer, K. D. Forbus, D. Gentner, The structure-mapping engine: Algorithm and examples, Artificial intelligence 41 (1989) 1–63.

[18] M. Kachroudi, Revisiting indirect ontology alignment: New challenging issues in cross-lingual context, arXiv preprint arXiv:2104.01628 (2021).

[19] Agno Team, Agno: Full-stack framework for building multi-agent systems, https://github.com/agno-agi/agno, 2025.

[20] Ontology Engineering Group (OEG), S4watr ontology documentation, https://oeg-upm.github.io/S4WATR/OnToology/S4Water.owl/documentation/index-en.html, 2024. Accessed: 2025-07-07.

[21] G. Carletti, E. Giulianelli, A. S. Lippolis, G. Lodi, A. G. Nuzzolese, M. Picone, G. Settanta, The water health open knowledge graph, arXiv preprint arXiv:2305.11051 (2023).

[22] F. Ahmedi, L. Ahmedi, B. O'Flynn, A. Kurti, S. Tahirsylaj, E. Bytyçi, B. Sejdiu, A. Salihu, Inwatersense: An intelligent wireless sensor network for monitoring surface water quality to a river in kosovo, International Journal of Agricultural and Environmental Information Systems (IJAEIS) 9 (2018) 39–61.

[23] Internet of Water Coalition / Ontology Engineering Group, Internet of water glossary, https://github.com/internetofwater/Glossary, 2025. Accessed: 2025-07-07.

[24] Open Geospatial Consortium, OGC *WaterML 2 – Part 3: Surface Hydrology Features (HY_Features)*, OpenGIS® Standard 14-111r6, OGC, 2018. URL: https://docs.ogc.org/is/14-111r6/14-111r6.html, accessed: 2025-07-07.