Ontology Evolution in Invasion Biology Using Large Language Models: A Hybrid Approach*

Hrishikesh Jadhav[1], Tina Heger[2], Birgitta König-Ries[3] and Alsayed Algergawy[1,3]

[1]Chair of Data and Knowledge Engineering, University of Passau, Germany
[2]Leibniz Institute of Freshwater Ecology and Inland Fisheries (IGB), Berlin, Germany
[3]Institute for Computer Science, University of Jena, Jena, Germany

Abstract

In previous work we developed a core ontology for the invasion biology domain, the INBIO. In that version, we modeled a part of the domain by considering concepts contained in a set of important hypotheses in this domain identified by earlier work, without considering other available resources such as publications testing these hypotheses. A next step is to update the INBIO ontology with respect to available resources. To this end, we propose a hybrid approach for ontology evolution that integrates Large Language Models (LLMs)—specifically GPT-4-based pipelines—with classical ontology engineering practices. This integration aims to create dynamic, scalable, and semantically consistent ontologies suitable for representing emergent phenomena in invasion biology. In particular, the proposed approach has three main components: extraction of concept and relationship candidates by analyzing hypothesis texts, scholarly abstracts, and curated domain metadata; usage of an LLM-driven pipeline (incorporating prompt-engineering and zero-shot learning) to generate novel concepts and relationships, linking previously unconnected ecological and socioeconomic attributes; and finally validation of newly proposed classes by domain experts in an iterative loop.

Keywords

Ontology Evolution, Semantic Web, Knowledge Graph, Large Language Models, Invasion Biology

1. Introduction

Ontologies, as the core component of the semantic web, facilitate data sharing, integration, and analysis by providing structured, machine-readable representations of domain knowledge [1, 2]. Ontologies encode knowledge of a specific domain in a well-structured representation, defining relationships between concepts and entities of the specified domain [1, 2]. As data generation in a domain continues to grow and adapt, there is the need to maintain ontologies up to date with respect to changes in the domain that they represent [3, 4]. Ontology evolution refers to the dynamic process of modifying and updating ontologies to reflect changes in knowledge, user requirements, or the environment, ensuring that the information remains accurate and relevant.

Invasion biology is concerned with the question why some species are able to establish and spread in an area where they have not evolved. Over time, the research community has developed several major hypotheses, and empirical studies have been conducted to test them [5]. In a previous work, we designed and developed a core ontology, called `Invasion Biology Ontology (INBIO)`[1], to represent knowledge from the invasion biology domain [6]. In the realm of invasion biology, the *INBIO* has sought to formalize fundamental concepts such as *enemy release*, *biotic resistance*, and *propagule pressure*. However, as novel findings accumulate, static ontologies often fail to incorporate emergent theories and observational data. Consequently, they risk obsolescence if not updated to reflect current scientific knowledge. For the development of `INBIO`, we relied on a compiled set of hypotheses that

---

LLM-TEXT2KG 2025: 4th International Workshop on LLM-Integrated Knowledge Graph Generation from Text

*Corresponding author.

jadhav02@ads.uni-passau.de (Hrishikesh Jadhav); t.heger@tum.de (T. Heger); birgitta.koenig-ries@uni-jena.de (B. König-Ries); alsayed.algergawy@uni-passau.de (A. Algergawy)
0000-0003-2763-8795 (Hrishikesh Jadhav); 0000-0002-5522-5632 (T. Heger); 0000-0002-2382-9722 (B. König-Ries); 0000-0002-8550-4720 (A. Algergawy)

[1] https://bioportal.bioontology.org/ontologies/INBIO

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

reflected the scientific understanding of invasion biology at a particular point in time. These hypotheses served as a foundational abstraction of domain knowledge. However, we did not incorporate other valuable resources available within the domain—such as publications on empirical studies conducted to test those hypotheses—which are essential for capturing the dynamic nature of scientific progress and theory evolution.

Ontology evolution is not an atomic process, but it consists of a number of different tasks. This includes identifying the need for evolution, determining the required changes, implementing these changes, and validating and assessing the new changes [3]. A number of approaches have been proposed to deal with these different tasks [3, 4, 7, 8]. However, most of these approaches lack the integration of these tasks in an intelligent environment. Ontology evolution is driven by continuous scientific developments, empirical discoveries, evolving terminology, and refinement of existing theories, requiring systematic updates to maintain semantic consistency and domain relevance.

Large Language Models (LLMs) offer an alternative to purely manual approaches by parsing unstructured text and suggesting structured outputs. In particular, transformer-based models like GPT-4 can propose new terms or relationships that domain experts might overlook [9]. Yet, relying solely on automated approaches can lead to inaccuracies or semantic drift, especially when context-specific or domain-specific meaning differs from common usage. Hence, there is a growing need for *hybrid* ontology evolution strategies that combine automated extraction with human validation [10].

This work aims to develop and validate a pipeline tailored to the rapidly evolving knowledge base of invasion biology. Specifically, we target three main objectives. First, we aim for scalability and timeliness by deploying an LLM-driven workflow that can quickly detect and integrate novel concepts or relationships, thereby minimizing manual curation burdens. Second, we focus on maintaining semantic alignment by cross-referencing new ontology elements with external repositories (e.g., BioPortal [11]) and employing automated reasoners within tools such as Protegé. Finally, we emphasize a balanced hybrid approach that combines data-driven suggestions from LLMs with expert oversight, ensuring that updates remain firmly grounded in scientific context.

2. Related Work

In this section, we outline the background and necessary preliminaries to the context of ontology evolution in the invasion biology domain using LLMs.

2.1. Invasion Biology and the HoH Framework

Invasion biology aims to elucidate factors driving non-native species success. The *Hierarchy-of-Hypotheses (HoH)* framework [12] organizes broad theoretical constructs—such as *enemy release* or *invasional meltdown*—into structured, testable propositions. INBIO provides a formal structure for these hypotheses by representing key ecological concepts and their interrelations, which can support the alignment of theoretical models. Yet, ongoing research yields new or revised statements on how ecological interactions or environmental gradients affect invasive processes. Traditional ontology curation efforts cannot always keep pace with these emergent insights, risking a gap between theory and knowledge representation.

2.2. Traditional Ontology Engineering and Evolution

Ontology engineering typically involves iterative steps: defining domain scope, enumerating crucial concepts, creating relationships, and refining constraints [13]. Frameworks like Methontology and NeOn provided structured methodologies for distributed ontology development [14]. However, when applied to highly dynamic fields—from invasion biology to cyber threat intelligence—manual ontology updates become unsustainable, risking outdated or inconsistent conceptual structures [15].

Ontology evolution specifically targets the continuous updates needed to keep an ontology aligned with new or revised domain knowledge [14]. Early studies emphasized detecting conceptual mismatches and concept drift, often relying on manual conflict resolution [16]. Although such methods ensured high precision, they struggled with scalability when domain knowledge changed frequently or in large increments [14].

2.3. Ontology Learning and Machine Learning

Research on *ontology learning* attempted partial automation of ontology engineering by leveraging text mining, pattern recognition, and clustering to extract new terms and their relations from large corpora [17, 18, 19]. Despite easing the manual burden, these systems often relied on shallow linguistic cues, leading to errors in polysemous terms [17].

The advent of transformer-based neural language models, including BERT, RoBERTa, and GPT variants, introduced deeper contextual understanding. They are capable of recognizing linguistic nuances that simpler NLP pipelines might miss. For instance, *REBEL* [20] demonstrate advanced extraction of domain-specific relationships from unstructured text. Nevertheless, purely automated pipelines risk introducing semantic conflicts or redundant concepts when domain oversight is minimal [21].

2.4. LLMs for Domain Adaptation

Transformer-based *Large Language Models (LLMs)* have broadened the possibilities for automatic knowledge extraction [9]. *Zero-shot and few-shot* learning techniques enable these models to interpret domain-specific corpora without extensive labeled training data [22]. *Prompt engineering* further steers model outputs, minimizing irrelevant or off-topic generation [23]. In parallel, *expert-in-the-loop* paradigms mitigate misclassifications, ensuring that domain context remains central to the final ontology updates [24].

While LLMs have been explored for tasks like summarizing scientific literature or detecting conceptual relations, relatively few studies integrate them into a full *ontology evolution* cycle. Challenges include preserving the ontology's logical consistency, preventing duplicate classes, and handling ambiguous terms that might have multiple senses across ecological, socioeconomic, or cultural contexts.

2.5. Research Gaps and Opportunities

To sum up, despite notable progress in ontology development and evolution, a number of issues and challenges require more attention. Many ontology learning systems ingest data on a one-time basis, lacking incremental or real-time update capabilities [10]. Changes to an ontology can propagate inconsistencies to linked datasets or reliant applications, requiring robust version control. Aligning newly generated terms with established ontologies (BioPortal, OBO Foundry, etc.) is often ad hoc, underscoring the need for systematic cross-referencing [25]. Additionally, although LLMs exhibit strong language understanding, they can hallucinate or produce domain-inaccurate definitions when specialized context is missing.

3. Methodology

To address these challenges, we propose a hybrid pipeline that combines LLM-facilitated concept discovery with domain-expert review and semantic validation, aiming to address these research gaps systematically. In this section, we explain the step-by-step methodology for data ingestion, LLM-driven extraction, expert validation, and ontology integration in the context of invasion biology.

3.1. Overall Pipeline

As we mentioned before, ontology evolution is not an atomic process, it consists of a number of overlapping and interconnected tasks. It starts by identifying the reasons behind the need to change and evolve the current ontology version. Then, parts that need to be changed are discovered. After that it applies these changes and finally these changes have to be validated and revised. Aligned to this scheme, we propose an ontology evolution pipeline composed of four key stages, as shown in Figure 1. To cover main tasks of ontology evolution, we propose a *data processing* step to check the need for evolution. This is carried out by analyzing available resources within the invasion biology domain. After that to check needed changes in terms of new concepts and relations, we introduce an *LLM-driven concept extraction* step. The *ontology update and expert validation* step is used to implement and validate the new changes. Finally, to assess these new changes, we offer a *competency question and reasoning* step. In the following, we are going to introduce detailed description of each step.

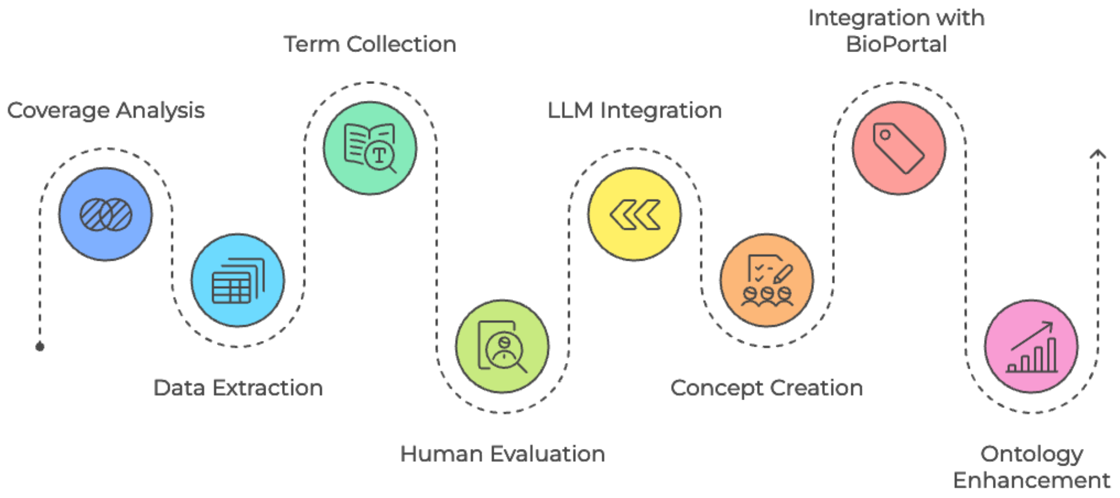

Figure 1: Overview of the ontology evolution pipeline. The process is organized into four main stages: **(1) Data Processing** — Coverage Analysis and Data Extraction; **(2) LLM-driven Concept Extraction** — Term Collection and LLM Integration; **(3) Ontology Update and Expert Validation** — Human Evaluation and Concept Creation; **(4) Competency and Reasoning Validation** — Integration with BioPortal and Ontology Enhancement.

3.2. Data Collection and Preprocessing

Data Sources. We gather data from three primary sources:

- **Invasion Biology Ontology (INBIO)**: A structured ontology covering key hypotheses in invasion biology, including *enemy release*, *biotic resistance*, and *invasional meltdown* [6]. The INBIO ontology provides a hierarchical framework for hypothesis representation and validation.
- **Curated Hypothesis Files**: A structured collection of domain hypotheses, manually extracted from research literature and stored in **Excel format**. This dataset includes hypotheses such as *Tens Rule* and *Propagule Pressure*, which are essential for understanding species establishment and spread. Each hypothesis file consists of structured metadata, key explanatory variables, and supporting evidence [6].
- **Research Abstracts**: A collection of peer-reviewed papers and scientific abstracts sourced from literature repositories. These abstracts provide emerging concepts, new invasion patterns, and domain-specific terminology that may not yet be reflected in structured ontologies.

Preprocessing Steps. The raw data typically contain duplicates, inconsistent encodings, and varying naming conventions for species and hypotheses. To address this:

- **Cleaning**: Remove duplicates and harmonize domain synonyms (e.g., *alien species* vs. *non-native species*).
- **Normalization**: Convert multi-lingual or region-specific terms into a unified format where feasible.

- **Transformation**: Tokenize textual data into sentences or paragraphs for efficient ingestion by the Large Language Model (LLM).

We employ Python scripts for batch processing and consistency checks, ensuring each text snippet is appropriately encoded and free of irrelevant noise.

3.3. LLM-Driven Concept Extraction

Model and Prompt Strategy. We adopt GPT-4 (via OpenAI's API) for parsing domain corpora and generating candidate classes, properties, or relationships. The prompt strategy includes:

1. **Contextual Few-Shot Examples**: Illustrations of known relationships, such as *"invasive plant"* → *"displaces"* → *"local vegetation"*.
2. **Structured Output**: Requested in a JSON-like format using Pydantic [11], capturing Subject, Predicate, Object, and optional definitions or synonyms.
3. **Zero-Shot Handling**: If no annotated samples exist, the model attempts domain inference from general knowledge and flags results for expert review.

Validation of Outputs. Concept extraction involves automated LLM processes with manual validation by domain experts to ensure accuracy and relevance. The pipeline specifically:

- Filters out off-topic suggestions (e.g., *ecosystem* → *has biotic factor*).
- Flags ambiguous terms for domain-expert scrutiny (e.g., *"resistance"* which could indicate ecological or socio-political forms).

Domain experts participate actively in reviewing flagged terms to ensure accurate and meaningful integration into the ontology.

3.4. Ontology Update and Expert Validation

BioPortal Cross-Referencing. Newly proposed classes (e.g., *"generalist herbivore"*) are first checked against BioPortal [11] to ensure alignment with established ontologies. This step detects synonyms, definitions, and potential hierarchical placements. Where no direct match is found, the term is treated as novel and assigned a provisional status. Ontology updates are considered valid if they accurately reflect current scientific consensus, resolve semantic ambiguities, maintain internal consistency, and effectively support domain-specific competency queries. Conflicts or ambiguous terms are resolved through iterative expert validation and consensus-building.

Expert-in-the-Loop Review. Each suggested concept or relationship undergoes domain validation by an expert panel. They categorize proposals into:

- **Accepted**: Semantically valid and novel, or a useful refinement.
- **Revised**: Minor corrections needed (e.g., merging with an existing property).
- **Rejected**: Irrelevant or incorrectly inferred by the LLM.

Accepted or revised elements are integrated into the updated INBIO ontology.

3.5. Technology Implementation

Our implementation uses an integrated set of technologies for effective ontology evolution. The implementation details can be found in our GitHub repository [2].

---

[2] https://github.com/EcoWeaver/DomainOntologies/tree/development

3.5.1. System Architecture

The system architecture combines artificial intelligence techniques with traditional ontology engineering practices:

- **Knowledge Base**: The INBIO core ontology is stored in OWL format using Protégé 5.5.0.
- **LLM Access**: GPT-4 integration via OpenAI API through Python (with the ‘openai‘ library).
- **Validation Tools**: Web-based BioPortal REST API queries.
- **Persistence**: Neo4j graph database for efficient storage of the knowledge graph.

3.5.2. LLM Integration and Prompting

Our approach employs specific prompt templates for ontology element extraction:

Concept Extraction Prompt:

Listing 1: Concept Extraction Prompt Template

```
Generate a detailed concept for the term '{term}'
in the context of invasion biology, incorporating
the following inputs:
- Hypothesis Text: '{hypothesis_text}'
- Abstract: '{abstract}'
- Ontology Schema (VERSION 1 INBIO): '{ontology_schema}'

Please provide information for each field:
1. Label: A clear name for the concept.
2. Definition: A precise explanation of the term.
3. Annotations: Additional context on significance.
4. Subclass Of: Identification of the parent concept.
```

This prompting helps filter system-generated noise and improves semantic alignment during the extraction phase.

3.5.3. BioPortal Integration

The BioPortal API is integrated to validate and enrich ontology terms. For example, when processing the term "invasive species", the system:

1. Constructs a query to BioPortal’s search endpoint 2. Extracts metadata like definitions, URIs, and related concepts 3. Incorporates this information into the validation workflow

This integration ensures consistency with broader biomedical vocabularies and reduces redundancy in the ontology.

4. Results

The proposed approach in this paper has been developed and implemented. More details are available at the GitHub repository[3]. Furthermore, we evaluate the proposed approach using datasets available at the HoH website[4].

4.1. Extraction Metrics

The LLM-driven ontology evolution pipeline systematically processed invasion biology literature, resulting in structured ontology elements categorized into distinct semantic types:

- **Subjects (69 unique terms)**: Identified as primary entities initiating actions or processes, such as *generalist herbivore* or *biotic acceptance.*

---

[3] https://github.com/EcoWeaver/DomainOntologies/tree/development

[4] https://www.hi-knowledge.org/tools

- **Predicates (63 unique terms)**: Relationships linking subjects to objects, including terms like *facilitates*, *influences*, or *disrupts*.
- **Objects (113 unique terms)**: Target entities involved in or impacted by actions, for example, *ecosystem dynamics*, *species distribution*, or *species coexistence*.

These semantic elements constituted **175 unique triplets**, explicitly structured as subject-predicate-object relations. The metrics of the extraction process, including expert validation outcomes, are summarized in Table 1.

Table 1
Summary of Extracted Terms and Triplets

Category	Count
Subjects	69
Predicates	63
Objects	113
Total Extracted Triplets	175
Expert-Validated Triplets for INBIO	83
Validation Accuracy (%)	47.43%

Interpretation of Metrics. The initial set of 175 candidate triplets generated by the LLM represented preliminary ontology suggestions. Each was manually validated by domain experts for semantic correctness, relevance, and ontology alignment. This rigorous assessment yielded 83 confirmed triplets, reflecting a validation accuracy of approximately 47.43%. This accuracy demonstrates both the pipeline's efficacy in generating meaningful suggestions and the importance of subsequent expert validation to maintain ontology quality.

Extraction Scope and Validation Procedure. Initially, approximately 1,200 terms were identified by the LLM from a broader textual corpus covering ecological, environmental, socioeconomic, and policy contexts. Expert validation was essential for identifying and removing irrelevant or redundant terms, resulting in the final selection of validated triplets that significantly contributed to the ontology's conceptual clarity and interdisciplinary scope.

4.2. Illustrative Triplets

Extracted triplets exemplify critical ecological interactions within invasion biology. Notable examples include:

- *invasion patterns* → *illustrate* → *colonisation patterns of species*
- *mutualistic interactions* → *facilitate* → *establishment of invasive species*
- *invasion* → *disrupts* → *indigenous ecosystems*

These examples underscore the pipeline's capacity to identify and articulate complex ecological phenomena.

4.3. Added Concepts and Relationships

The ontology evolution pipeline proposed 46 new concepts and 24 new relationships, with detailed statistics provided in Table 2.

Representative concept definitions and relationships include:

- **Generalist Herbivore**:

Table 2
Summary of Added Concepts and Relationships

Category	Count
New Concepts Introduced	46
New Relationships Added	24
Validated Concepts and Relationships	>90%

  - *Definition*: An organism feeding broadly across diverse plant species rather than specialized diets.
  - *Relationships*: Impacts ecosystem dynamics.
- **Biotic Acceptance**:
  - *Definition*: The phenomenon wherein ecosystems with high native biodiversity levels more readily integrate non-native species, potentially influencing overall species richness.
  - *Relationships*: Affects native biodiversity.

4.4. Expert Validation Results

The expert validation process critically evaluated the relevance, precision, and consistency of proposed ontology additions through iterative reviews by domain specialists in invasion biology. The procedure classified outputs into categories: Accepted Without Modifications (68%), Accepted With Modifications (22%), and Rejected (10%), as detailed in Table 3.

Table 3
Summary of Expert Validation Feedback

Feedback Type	Percentage
Accepted Without Modifications	68%
Accepted With Modifications	22%
Rejected	10%

Critical modifications derived from expert input included:

- Reclassification of entities (e.g., Alien and Native Species reclassified directly under *conceptual entity*).
- Removal of redundant and obsolete concepts such as *plant weeds*.
- Refinement of relationships (e.g., improved relationship representation for "generalist herbivore" as "affects ecosystem dynamics").

4.5. Competency Questions and SPARQL Queries

Competency questions (CQs) formulated by domain experts played a pivotal role in ontology evaluation, serving as structured benchmarks for ontology adequacy, consistency, and completeness. A total of five core competency questions were systematically evaluated using SPARQL queries against the updated ontology:

1. **What is an invasive species?**
2. **What is an alien species?**
3. **List all classes with their definitions.**
4. **What does invasion biology examine?**
5. **What ecological functions does a detritivore enhance?**

The following is an illustrative example SPARQL competency query (CQ1):

Listing 2: SPARQL Query: Definition of Invasive Species

```
SELECT ?definition WHERE {
    ?class rdfs:label "invasive species" .
    ?class skos:definition ?definition .
}
```

Query Result:

> “Alien species that sustain self-replacing populations over several life cycles, produce reproductive offspring, often in very large numbers at considerable distance from the parent and/or site of introduction, and have the potential to spread over long distances.”

This definition matches established scientific literature, demonstrating successful ontology alignment with established domain knowledge.

Expert Involvement in Query Validation. Domain experts systematically verified all SPARQL query responses, ensuring correctness of definitions, alignment with ecological theories, and appropriateness within invasion biology’s scientific discourse. This expert oversight validated the ontology’s robustness, highlighted subtle semantic discrepancies, and informed targeted modifications to enhance ontology reliability and relevance.

In summary, the LLM-driven ontology pipeline efficiently proposed substantial expansions, systematically validated through structured expert reviews and competency-based querying, thereby ensuring semantic accuracy, interdisciplinary comprehensiveness, and practical applicability to invasion biology research and applications.

5. Discussion

In this section, we discuss the pros and cons of the proposed approach as well as areas for future direction:

5.1. Advantages of the Hybrid Approach

Our methodology demonstrates several key benefits for ontology evolution in dynamic scientific domains:

- Accelerated Knowledge Acquisition. The LLM-based extraction significantly reduced the time required to identify candidate terms from new literature. What previously took 2-3 hours of manual review per paper was reduced to approximately 10 minutes of validation time, with the LLM handling the initial extraction in seconds.
- Novel Connection Identification. The GPT-4 model frequently suggested relationships between concepts that were not explicitly stated in the source text but were semantically valid. For example, it correctly associated the "Naturalization-Interference Paradox" with both "Competition Mechanisms" and "Succession Dynamics" despite these connections being implicit in the literature.
- Reduced Expert Fatigue.By pre-filtering suggestions and providing structured candidate terms, the approach substantially reduced cognitive load on domain experts, allowing them to focus on validation rather than comprehensive manual extraction.

5.2. Limitations and Challenges

Despite the promising results, several challenges remain:

- Domain-Specific Ambiguity. LLMs occasionally suggested terms that seemed valid from a general knowledge perspective but had different meanings in invasion biology. For example, "resistance" was sometimes incorrectly associated with antibiotic resistance rather than community resistance to invasion.
- Neologism Handling. Newly coined terms in recent literature presented challenges, as they lacked sufficient representation in the LLM's training data. Supplementing with specialized glossaries partially mitigated this issue.
- Hierarchical Placement. While the LLM excelled at identifying concepts, it was less reliable in suggesting optimal taxonomic placements within the ontology hierarchy. This aspect still required significant expert input.

5.3. Scalability and Dependency Management

Adopting a modular architecture supports batch-mode or near real-time integration. For instance, incremental updates can be triggered whenever a new dataset or publication hits a certain threshold of domain relevance. Although real-time updates carry a higher computing cost, they ensure INBIO remains updated with minimal lag, a feature critical for fast-moving fields.

The pipeline design can generalize to other data-intensive fields (e.g., epidemiology, climate change studies) where new findings emerge rapidly, requiring frequent ontology updates. The key is ensuring domain experts remain an integral part of the loop.

5.4. Future directions

Our work directly contributes to the incremental evolution of the INBIO ontology, effectively resulting in successive ontology versions. The validated concepts, properties, and relationships are systematically integrated into formal releases, marking explicit milestones in the ontology's lifecycle. Each validated update results in an explicit new version INBIO 1.2, reflecting structured improvements over earlier ontology snapshots. Beyond the addition of new concepts, future iterations should explicitly address ontology maintenance tasks including:

- **Updating Existing Concepts**: Regular revision of ontology entries is necessary, involving refinement of definitions, improvement of hierarchical placements, and resolution of semantic ambiguities as new domain insights emerge.
- **Deletion of Obsolete or Redundant Concepts**: Implement systematic processes to remove or merge outdated or redundant terms, thereby maintaining ontology coherence and clarity.

Additional promising avenues for future research and improvement include:

- Real-Time Ontology Updates. Developing automated mechanisms for continuous monitoring and incorporation of newly published research, allowing near-instantaneous ontology evolution.
- Cross-Ontology Alignment. Expanding interoperability by aligning INBIO with complementary ontologies such as the Environment Ontology (ENVO) and the Population and Community Ontology (PCO), facilitating cross-domain knowledge exchange.
- Multilingual Extensions. Adapting and extending the pipeline for multilingual data ingestion to incorporate international perspectives and enhance comprehensiveness across diverse ecological and policy contexts.
- Interactive Visualization. Creating tools to visually represent ontology evolution over successive versions, clearly depicting changes, additions, and shifts in thematic emphasis, thereby supporting domain-expert exploration and stakeholder communication.

6. Conclusion

In this paper, we introduced a hybrid methodology for ontology evolution that leverages Large Language Models to accelerate concept extraction while maintaining semantic rigor through expert validation. Using the Invasion Biology Ontology (INBIO) as our case study, we demonstrated significant expansion in domain coverage, particularly in previously underrepresented socioeconomic and management dimensions. The combination of automated extraction with expert-driven validation significantly improved the scalability and responsiveness of ontology updates, establishing a systematic approach for continuous ontology evolution. While our approach substantially reduces manual workload, it relies on manual hierarchical placement and expert review due to occasional domain-specific ambiguities. Additionally, LLM-driven extraction currently suffers from moderate accuracy (47%), indicating potential for improvement.

To address these limitations, we plan future work on comparative benchmarking with other ontology evolution methods, fine-tuning domain-specific language models, implementing advanced alignment tools (AML [26], LogMap [27]), and integrating multilingual capabilities to further improve accuracy and comprehensiveness.

In conclusion, integrating advanced language models and rigorous semantic web methodologies demonstrates substantial potential to manage dynamic and continuously expanding scientific knowledge domains. Continued enhancements in LLM capabilities, ontology engineering frameworks, and expert validation processes promise further advances in semantic knowledge representation.

Acknowledgments

This research was partially conducted within the context of the ZiF Resident Group "Mapping Evidence to Theory in Ecology" at the Center for Interdisciplinary Research (ZiF), Bielefeld University. We acknowledge the support and discussions within this group, which contributed to the refinement of our methodology.

This work was also supported by the German Research Foundation (DFG) through the INAS project, grant number 455913229.

Declaration on Generative AI

Generative Artificial Intelligence (AI) tools were utilized in the preparation of this manuscript. Specifically, the authors used generative AI to correct grammar, improve language clarity, and enhance the scientific writing style of paragraphs written initially by the authors. At no point were generative AI tools employed to generate original scientific content or to create paragraphs that were not first drafted by the authors. All ideas, scientific contributions, and research findings presented in this paper originate from the authors themselves.

References

[1] E. K. Jacob, Ontologies and the semantic web, Bulletin of the American Society for Information Science and Technology 29 (2003) 19–19.

[2] L. Ding, P. Kolari, Z. Ding, S. Avancha, Using ontologies in the semantic web: A survey, Ontologies: A Handbook of Principles, Concepts and Applications in Information Systems (2007) 79–113.

[3] F. Zablith, G. Antoniou, M. d'Aquin, G. Flouris, H. Kondylakis, E. Motta, D. Plexousakis, M. Sabou, Ontology evolution: a process-centric survey, The knowledge engineering review 30 (2015) 45–75.

[4] A. Groß, C. Pruski, E. Rahm, Evolution of biomedical ontologies and mappings: overview of recent approaches, Computational and structural biotechnology journal 14 (2016) 333–340.

[5] J. M. Jeschke, T. Heger, Invasion Biology: hypotheses and evidence, CAB International, 2018.

[6] A. Algergawy, R. Stangneth, T. Heger, J. M. Jeschke, B. König-Ries, Towards a core ontology for hierarchies of hypotheses in invasion biology, in: European Semantic Web Conference, Springer, 2020, pp. 3–8.

[7] P. Plessers, O. De Troyer, S. Casteleyn, Understanding ontology evolution: A change detection approach, Journal of Web Semantics 5 (2007) 39–49.

[8] H. Kondylakis, N. Papadakis, Evordf: Evolving the exploration of ontology evolution, The Knowledge Engineering Review 33 (2018) e12.

[9] F. Neuhaus, Ontologies in the era of large language models – a perspective, Applied Ontology 18 (2023) 399–407. doi:`10.3233/AO-230072`.

[10] H. Dong, J. Chen, Y. He, Y. Gao, I. Horrocks, A language model based framework for new concept placement in ontologies, 2024. `arXiv:2402.17897`.

[11] P. L. Whetzel, N. F. Noy, N. H. Shah, P. R. Alexander, C. Nyulas, T. Tudorache, M. A. Musen, Bioportal: Enhanced functionality via new web services from the national center for biomedical ontology to access and use ontologies in software applications, Nucleic Acids Research 39 (2011) W541–W545. doi:`10.1093/nar/gkr469`.

[12] T. Heger, C. A. Aguilar-Trigueros, I. Bartram, R. R. Braga, G. P. Dietl, M. Enders, D. J. Gibson, L. Gómez-Aparicio, P. Gras, K. Jax, S. Lokatis, C. J. Lortie, A.-C. Mupepele, S. Schindler, J. Starrfelt, A. D. Synodinos, J. M. Jeschke, The hierarchy-of-hypotheses approach: A synthesis method for enhancing theory development in ecology and evolution, BioScience 71 (2020) 337–349. doi:`10.1093/biosci/biaa130`.

[13] N. F. Noy, D. L. McGuinness, Ontology Development 101: A Guide To Creating Your First Ontology, Technical Report, Stanford University, Stanford, CA, 94305, 2001. Available at: `noy@smi.stanford.edu` and `dlm@ksl.stanford.edu`.

[14] F. Zablith, G. Antoniou, M. d'Aquin, G. Flouris, H. Kondylakis, E. Motta, D. Plexousakis, M. Sabou, Ontology evolution: a process-centric survey, The Knowledge Engineering Review 30 (2013) 45–75. doi:`10.1017/S0269888913000349`, first published online 28 August 2013.

[15] N. Noy, M. Klein, Ontology evolution: Not the same as schema evolution, Knowledge and Information Systems (2004). doi:`10.1007/s10115-003-0137-2`, source: DBLP.

[16] P. Plessers, O. De Troyer, S. Casteleyn, Understanding ontology evolution: A change detection approach, Journal of Web Semantics 5 (2007) 39–49. doi:`10.1016/j.websem.2006.11.001`.

[17] M. N. Asim, M. Wasim, M. U. G. Khan, W. Mahmood, H. M. Abbasi, A survey of ontology learning techniques and applications, Database 2018 (2018) bay101.

[18] H. B. Giglou, J. D'Souza, S. Auer, Llms4ol: Large language models for ontology learning, 2023. URL: https://arxiv.org/abs/2307.16648. `arXiv:2307.16648`.

[19] P. Mateiu, A. Groza, Ontology engineering with large language models, 2023. URL: https://arxiv.org/abs/2307.16699. `arXiv:2307.16699`.

[20] H. Babaei Giglou, J. D'Souza, S. Auer, LLMs4OL: Large Language Models For Ontology Learning, in: Proceedings of the Ontology Learning Workshop, TIB Leibniz Information Centre for Science and Technology, Hannover, Germany, 2023.

[21] Y. He, J. Chen, H. Dong, I. Horrocks, Exploring large language models for ontology alignment, 2023. URL: https://arxiv.org/abs/2309.07172. `arXiv:2309.07172`.

[22] Y. Li, A practical survey on zero-shot prompt design for in-context learning (2023). doi:`10.26615/978-954-452-092-2_069`.

[23] J. White, Q. Fu, S. Hays, M. Sandborn, C. Olea, H. Gilbert, A. Elnashar, J. Spencer-Smith, D. Schmidt, A prompt pattern catalog to enhance prompt engineering with chatgpt, ArXiv abs/2302.11382 (2023). doi:`10.48550/arXiv.2302.11382`.

[24] J. Walker, E. Koutsiana, M. Nwachukwu, A. Meroño Peñuela, E. Simperl, The promise and challenge of large language models for knowledge engineering: Insights from a hackathon, in: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, CHI EA '24, Association for Computing Machinery, New York, NY, USA, 2024. doi:`10.1145/3613905.3650844`.

[25] J. Raad, A. Bertaux, C. Cruz, A survey on how to cross-reference web information sources, in: Science and Information Conference (SAI), 2015, IEEE, 2015, pp. 609–618.

[26] 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, volume 8185 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2013, pp. 527–541.

[27] E. Jiménez-Ruiz, B. C. Grau, LogMap: Logic-Based and Scalable Ontology Matching, in: Proceedings of the 10th International Semantic Web Conference (ISWC), volume 7031 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 273–288.