

# Mining for Species, Locations, Habitats, and Ecosystems from Scientific Papers in Invasion Biology: A Large-Scale Exploratory Study with Large Language Models

Jennifer D’Souza<sup>1</sup>, Zachary Laubach<sup>2</sup>, Tarek Al Mustafa<sup>3</sup>, Sina Zarriëß<sup>4</sup>,  
Robert Frühstückl<sup>5</sup>, Phyllis Illari<sup>6</sup>

<sup>1</sup>TIB Leibniz Information Centre for Science and Technology,

<sup>2</sup>University of Colorado Boulder, <sup>3</sup>Friedrich Schiller University Jena, <sup>4,5</sup>Bielefeld University,

<sup>6</sup>University College London

jennifer.dsouza@tib.eu

## Abstract

This paper presents an exploratory study that harnesses the capabilities of large language models (LLMs) to mine key ecological entities from invasion biology literature. Specifically, we focus on extracting species names, their locations, associated habitats, and ecosystems, information that is critical for understanding species spread, predicting future invasions, and informing conservation efforts. Traditional text mining approaches often struggle with the complexity of ecological terminology and the subtle linguistic patterns found in these texts. By applying general-purpose LLMs without domain-specific fine-tuning, we uncover both the promise and limitations of using these models for ecological entity extraction. In doing so, this study lays the groundwork for more advanced, automated knowledge extraction tools that can aid researchers and practitioners in understanding and managing biological invasions.

## 1 Introduction

Human population growth and expansion are coupled with the intentional and unintentional movement of species beyond their historic ranges. The introduction of nonnative species can have dramatic impacts that cascade across ecological scales from individual plants and animals to populations and communities (Roy et al., 2023). The goal of invasion biology is to identify and understand the impacts of alien species on native flora and fauna across these ecological scales, not only to conserve rare, sensitive, and ecologically valuable native species, but also because intact functional ecosystems provide important economic, public health, and socioemotional services to humans (Casseo et al., 2018; Jeschke and Heger,

2018). Achieving this goal becomes increasingly challenging though, since alien species introductions operate at a rapid rate and global scale, and in the context of increasing human population growth and climate change. Additionally, the body of research in this domain of ecology is currently growing to an extent that it becomes more and more difficult for invasion biology researchers to build a systematic and shared understanding of the impacts of alien species on ecosystems, or to even comprehensively categorize these species, and their locations, and relationships. This paper explores the potential of recent NLP technologies, such as Information Extraction (IE) approaches based on Large Language Models (LLMs) (Amatriain et al., 2023; D’Souza, 2023), as a tool that could contribute to predicting future invasions and their consequences.

The extraction and categorization of information from scientific publications is a well-known NLP task (Augenstein et al., 2017; Gábor et al., 2018; Luan et al., 2018; Brack et al., 2020; Dessì et al., 2020; D’Souza et al., 2021; Liu et al., 2021; Kabongo et al., 2021; D’Souza and Auer, 2022; D’Souza, 2024; Shamsabadi et al., 2024; D’Souza et al., 2024). In the biomedical domain, NER and RE facilitate large-scale biomedical data analysis, such as network biology (Zhou et al., 2014), gene prioritization (Aerts et al., 2006), drug repositioning (Wang and Zhang, 2013) and the creation of curated databases (Li et al., 2015). In the clinical domain, NER and RE can aid in disease and treatment prediction, readmission prediction, deidentification, and patient cohort identification (Miotto et al., 2018). However, the domain of ecology and invasion biology in particular has hardly been explored and dedicated datasets are scarce. To the best of our knowledge, the small-scale INAS dataset (Brinner et al., 2022) presents the only invasion biology-specific resource that provides annotations of hypotheses for scientific abstracts.

This paper presents a study on information extraction (IE) in invasion biology, encompassing both named entity recognition (NER) and relation extraction (RE). We simultaneously build on studies showing that jointly learning NER and RE can enhance overall performance (Giorgi et al., 2019) and on recent LLMs which may open new opportunities for IE. Thus, our central question is whether LLMs, with their advanced pattern recognition capabilities, can be effectively applied to a new domain to simultaneously identify entities and infer their relationships. We present an approach that prompts LLMs to identify the following four key entities in invasion biology—viz. species, location, habitat, and ecosystem. Since we cannot evaluate the results quantitatively, we explore them qualitatively with respect to the following questions: (i) do the LLMs extract reasonable species, location, habitat, and ecosystem information and what interactions do they infer? (ii) what relation information is mined?, (iii) what are the benefits of defining LLM workflows to mine large amounts of data? To summarize, this work makes two key contributions: (i) the release of a text data mining corpus of over 10,000 invasion biology papers, including titles, abstracts, full text for nearly 2,000 papers, and structured information extracted by the GPT-4o LLM from the abstracts, to support further research, analysis, and curation (<https://doi.org/10.5281/zenodo.13956882>); and (ii) a systematic workflow for schema discovery in information extraction tasks, demonstrated in this paper but broadly applicable for leveraging LLMs in open-ended IE objectives.

## 2 Our Text Data Mining Corpus

Before tackling the IE task, we compiled a publication corpus as the unstructured source for scientific information. This section outlines the preparation of this dataset for text mining and IE.

### 2.1 Defining the Collection

We started with the Invasion Biology Corpus (Mitchen et al., 2024), which lists metadata for 49,438 invasion biology papers in Wikidata. Using the papers’ DOIs, we queried the ORKG ASK search engine’s API to retrieve abstracts and full texts. ASK was chosen for its convenient access to over eight million publications (Knoth et al., 2023) across diverse scientific fields.

### 2.2 Corpus Statistics

Of the 49,438 DOIs queried, 12,636 were available in the ASK database. Among these, 9,802 had abstracts but not full text, and 2,834 included both. This distribution reflects a common issue: open-access full-text articles are limited, posing a challenge for scientific NLP efforts. As a result, our final dataset—consisting of a larger set of abstracts and a smaller set of full texts—serves as the target collection for text mining and IE.

The abstracts range from 10 to 1,608 tokens (average: 235), while the full texts span 28 to 123,958 tokens (average: 7,667), highlighting the greater informational depth of full texts despite their limited availability.

### 2.3 Bibliometric Analysis

Before the IE tasks, we highlight two key bibliometric insights from our corpus. Among 12,636 papers with abstracts, publication years span 52 years, starting in 1950. Full-text availability begins in 1990. Figure 1 shows a snapshot of the past 20 years, with 2016 having the most abstracts (1,183) and 2017 the most full texts (294). Figure 2 highlights the top ten publishers. Additional insights are available in our online repository.




Figure 1: Distribution of papers in our corpus with abstracts and with full text over the past 20 years.

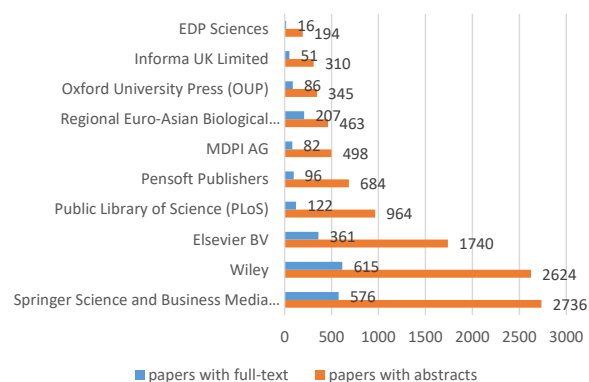


Figure 2: Distribution of papers, in our corpus, by abstract and full-text availability across the top ten publishers.

### 3 Information Extraction (IE) with Large Language Models (LLMs)

An IE task entails two prerequisites. 1) The collection of papers on the which the IE task should be performed. And 2) the schema which defines the IE extraction targets of interest. Having compiled our publication corpus above, we proceed with IE.

#### 3.1 Schema Discovery

Schema or semantic model discovery is a key focus of this subsection. We aim to use LLMs to propose a standard semantic structure for extracting information from each paper. This standardized IE format facilitates downstream processing of the structured data. Since no predefined set of relations exists, the schema must flexibly capture extracted entities and represent their relationships.

We approach the task in two main stages: **specialize** and **generalize**. In the **specialize** stage, the LLM is instructed to generate a schema as the target for IE, focusing on four entities: *species*, *location*, *habitat*, and *ecosystem*. The schema proposed by the LLM includes an IE objective with properties that specify relationships between the extracted entities. This behavior is guided by a system prompt, and the LLM processes each paper instance individually through user prompts to define a semantic model tailored or specialized to each paper. In the **generalize** stage, the LLM is instructed to generate a generalized schema based on input of the specialized schemas created per paper in the first stage. The generalized schema would then represent the most flexible way to capture the relations between the entities across all papers. The following sections provide detailed descriptions of each stage.

##### 3.1.1 Stage Specialize: Schemas per Paper

The LLM is the central tool in this endeavor, operating in the text *completion* or *generation* mode. The interaction with the LLM is guided by two types of prompts: the **SYSTEM PROMPT**, which defines the model’s persona and specialized behavior for a given task, and the **USER PROMPT**, which supplies input data for the defined behavior. The LLM’s behavior, once established by the system prompt, remains consistent throughout the interaction session per stage.

**SYSTEM PROMPT.** The system prompt defines the LLM’s persona and provides structured guidance for the task. It is composed of three key

components: (a) Role, (b) Task instruction, and (c) Output format. In this stage, the LLM assumes the overarching role of a “research assistant in invasion biology or ecology tasked with reading and understanding scientific papers for extracting relevant information within a schema.” Its primary instruction is to extract terms related to four entities—*species*, *location*, *habitat*, and *ecosystem*—and identify the relationships expressed in each paper. Additionally, the LLM is directed to structure the extracted information into an external container, representing the schema that defines the extraction target.

To refine the LLM’s performance, the system prompt was iteratively improved over two runs. In the first run, the prompt lacked detailed definitions of the entities. After discussions with an ecologist, the second run included precise definitions of the field of invasion biology and the four entities (see Table 1). The final system prompt, incorporates these definitions, enhancing the LLM’s clarity and consistency in identifying and extracting the specified entities and their relationships.

This refinement aligns with the principles of in-context learning (Radford et al., 2019), which emphasize that providing clear and detailed task instructions ensures the model’s better comprehension of the downstream task, leading to more accurate and consistent results.<sup>1</sup> By supplying the LLM with a well-defined context, we enable it to successfully extract and organize the desired information from scientific texts.

**USER PROMPT.** The model is given one paper at a time and is allowed to suggest its own semantic model. The instruction in the user prompt is simple and is as follows: “Extract the information as instructed from this article title and abstract. \n\n **Title:** {title}\n **Abstract:** {abstract}\n”.

**Result.** Ten papers were randomly selected, and the LLM was prompted to generate a schema for each, which was stored and is available in the data folder of our online repository. The tenth paper, an outlier from the Wikidata Invasion Biology dataset, highlights the potential presence of false positives—papers unrelated to invasion biology.

The schemas evolve significantly to better capture relationships across the nine true-positive pa-

---

<sup>1</sup>We acknowledge extensive literature on in-context learning advocating for few-shot task examples (Brown et al., 2020). In this paper, we interpret in-context learning as only providing a detailed task description without examples.

Entity	Description
<b>Species</b>	Includes specific named species (e.g., <i>Asterias amurensis</i> ) and broader categories (e.g., demersal fish, aquatic invertebrates), covering plants, animals, fungi, or microbes introduced to new environments where they establish, spread, and cause ecological or economic impacts. Higher-level taxonomic or functional groups are included when specific species are not identified, but generic terms like “invasive species” are excluded.
<b>Location</b>	Refers to study sites, from specific locations (e.g., “Port Phillip Bay, southern Australia”) to broader regions (e.g., southern Australia, Amazon rainforest). Includes natural features (rivers, bays, mountains) and administrative areas (cities, states, countries).
<b>Ecosystem</b>	A system of interacting biological and abiotic components, often spanning multiple locations (e.g., the savannah ecosystem across Kenya and Tanzania).
<b>Habitat</b>	A specific part of an ecosystem where an organism lives, such as crocodiles in freshwater habitats (e.g., rivers) within the savannah ecosystem.

Table 1: Definitions of the four entities that encompass the information extraction (IE) aim of this paper.

pers. Early schemas, like Schema 1, use basic categorizations of species, locations, and ecosystems with simple binary relationships. In contrast, later schemas, such as Schema 8 and Schema 9, adopt more sophisticated structures, accommodating complex interactions involving multiple entities and contextual factors. E.g., relationships like “introduced\_species\_location” or “ecological\_disturbance” incorporate broader ecological and anthropogenic processes, moving beyond pairwise connections. This progression enhances granularity, capturing nuanced interactions like species dynamics influenced by environmental conditions or human activity, making the schemas more meaningful for analysis and synthesis.

The schemas reveal recurring patterns and unique adaptations in representing ecological data. Common themes, such as invasion biology, pollination networks, and anthropogenic impacts, highlight key areas of focus. Standardized fields like *species* and *location* ensure consistency, while tailored relationship labels, such as “most effective pollinators” (Schema 2) and “competitive replacement” (Schema 5), add specificity and contextual relevance. The inclusion of spatial and environmental parameters, like *habitats* and *ecosystems*, emphasizes their role in shaping ecological interactions. These schemas balance standardization and flexibility, illustrating how modular design can synthesize diverse unstructured text data while preserving study-specific distinctions.

### 3.1.2 Stage Generalize: Generic Schema

In this stage, the goal was to develop a standardized container for the extracted information in the form of a JSON data structure, specifically de-

signed to capture the relationships between four entities: species, location, habitat, and ecosystem. The system prompt given to the LLM was similarly structured as the **specialize** stage, with the model’s role defined not only as a research assistant in invasion biology but also as having expertise in semantic modeling. Overall, the prompt consisted of the same structural components: (a) Role, (b) Task instruction, and (c) Output format. However, aside from the role, the task instruction defining the LLM’s behavior in this stage was different. At this stage, the LLM was expected to review each individual schema from the earlier stage for every paper and then thoughtfully propose a standardized schema. This is related to prior work (Baazizi et al., 2017, 2020) which presents an approach based on discovering a schema for individual JSON documents and then merging each of these schemas. In our case, we discover schemas from the unstructured texts of individual papers and then merge them in this stage. This schema would eventually form the IE objective, to be applied across all papers in the collection to extract the relevant entities and their relationships.

Since LLMs are known to generate different responses to the same prompt, the LLM was prompted thrice with the same task each time prompting the model in the same way: “Read the nine different schema instances and generate a standardized schema in the JSON output format.”

**Result.** The three JSON schema variants generated represented a thoughtful approach to standardizing the representation of entities and their relationships in invasion biology. The common thread across the 3 standardized schema responses

Extraction Target	Extracted Item	Extracted Item Properties
<b>Species</b>	name: species_name	role: native/introduced/alien/invasive taxonomy_level: species/genus/family
<b>Location</b>	name: location_name	category: natural/administrative geopolitical_info: country/region/city additional_details: climatic/physiographic
<b>Ecosystem</b>	name: ecosystem_name	type: aquatic/terrestrial/marine scope: local/regional/global
<b>Habitat</b>	name: habitat_name	type: aquatic/terrestrial/marine subcomponent_of: ecosystem_name specifics: e.g., benthic, littoral
<b>Relationships</b>	related_entities: [entity1, entity2, ...]	name: relationship_name type: biological/physical/ecological/anthropogenic directionality: unidirectional/bidirectional context: relationship_contextual_description

Table 2: Standardized information extraction (IE) schema for four ecological entities, their relationships, and associated properties, pertinent to structure information from invasion biology scientific papers.

was designed to capture information about four primary entities—species, location, habitat, and ecosystem—along with their properties and inter-relations, enabling consistent and structured extraction of relevant data. The standardized schema response 1 emphasized geospatial and hierarchical context, incorporating coordinates for locations and explicitly linking habitats to ecosystems. Schema 2 introduced detailed descriptions and a focus on the roles of species (e.g., native, introduced, invasive), while expanding relationship types to include biological, physical, and anthropogenic interactions. The standardized schema response 3 provided rich categorization, including species taxonomy and physiographic attributes for locations, along with specific habitat details such as benthic and littoral zones. Together, these schemas, while sharing commonalities, reflected minor differences in terms of the respective knowledge capture priorities, from geospatial precision to descriptive and taxonomic richness.

Based on these observations, we finalized the schema, which organizes data around core entities—species, locations, ecosystems, habitats, and relationships—each with structured properties tailored to ecological contexts. For example, species are characterized by roles (e.g., native, invasive) and taxonomy levels (e.g., genus, family), while locations are enriched with geopolitical and physiographic details. Ecosystems and habitats are distinguished by type and scope, with hierarchical

links between habitats and ecosystems enhancing clarity. Relationships capture interactions between entities, defined by type (e.g., biological, ecological) and directionality. This schema enables precise mapping of ecological networks and differentiation of key factors, such as invasive versus native species, across datasets. A detailed breakdown of the schema is presented in Table 2.

### 3.2 Information Extraction

With a standardized semantic structure for extracting information from each paper, enabling easier downstream processing, the LLM-based IE task was conducted.

#### 3.2.1 Stage Extract: Populate Schema

This stage now fulfills the main objective of this work, i.e. to extract information from a large-scale corpus (12,636 in our case) with an LLM to mine species, location, habitat, and ecosystem entities and their relations. The system prompt in this stage was close to the **specialize** stage system prompt where the role specified for the LLM was “research assistant in invasion biology or ecology tasked with reading and understanding scientific papers to extract relevant information *per the given predefined schema*.”

### 3.3 Technical Details

The proprietary OpenAI GPT-4o model was used for all tasks in this paper. Schema generation tasks in the **specialize** (Section 3.1.1) and **generalize**

(Section 3.1.2) stages were completed in a few seconds per schema. For the full extraction task in the **extract** stage (Section 3.2.1), which applied the extraction schema to a corpus of 12,636 papers, the LLM took approximately three days.

The resulting dataset is available on Zenodo. The code for running the GPT-4o model is accessible in our GitHub repository and supports processing either a single paper or a collection of papers.

### 3.4 Results and Discussion

Of the 12,636 papers in our dataset, the LLM classified 1,740 as outside the scope of invasion biology, responding with “N/A” and skipping extraction for these papers. This left 10,896 papers as the target for IE. This section summarizes the results.

The first exploratory insight highlights the diverse roles of species extracted by the LLM, reflecting their origins, behaviors, ecological functions, and impacts within invasion biology. Broad categories include **native**, **alien**, **introduced**, **invasive**, and **naturalized**, alongside more specific roles such as **agricultural weeds**, **biological control agents**, **pathogens**, **mutualists**, and **ecosystem engineers**. Some roles emphasize species’ origins (e.g., **indigenous**, **non-native**, **cryptogenic**), while others highlight their behaviors (e.g., **invasive**, **colonizer**, **expanding**) or ecological functions (e.g., **symbiont**, **facilitator**, **pioneer**). Additionally, certain roles capture species’ interactions within ecosystems, such as **co-introduced species**, **specialist herbivores**, or **cryptic invaders**, while others underline their relevance to conservation and management, such as **natural enemies**, **candidate biological control agents**, or **quarantine pests**. Together, this spectrum of roles underscores the complexity of species dynamics in invasion biology, providing valuable insights into biodiversity patterns, ecosystem impacts, and strategies for conservation and management. A detailed list of all roles is available here.

We conducted a finer-grained exploration of the extracted data, focusing on the most prevalent species names within the roles of *invasive*, *native*, and *introduced* species. **Invasive species** dominated, with examples such as *Procambarus clarkii* (76 mentions), *Harmonia axyridis* (73 mentions), and *Rhinella marina* (68 mentions) highlighting well-documented invaders. **Native species**, though less frequent, included specific examples like *Austroptamobius pallipes* and *Phragmites*

*australis* (24 mentions each). **Introduced species**, while fewer, included *Oncorhynchus mykiss* and *Crassostrea gigas*, reflecting their dual perspectives of their ecological integration and potential invasiveness. Overall, the data underscores the breadth of roles species play in invasion biology, from invaders disrupting ecosystems to native species requiring conservation focus and introduced species with varying impacts. However, the extraction also included generic terms (e.g., “native species” and “native plants”) as species names, indicating noise from the unsupervised nature of the IE task and highlighting the need for post-filtering to refine the analysis. The full list is available here.

The dataset highlights key **geopolitical locations**, with the five most frequent countries being Australia (406), South Africa (248), New Zealand (236), Italy (187), and France (168). Prominent regions include Europe (601), North America (348), the Mediterranean Sea (117), Asia (112), South America (98), the Mediterranean (80), and the Iberian Peninsula (75). Less common mentions were cities like Sydney (8), Hong Kong (7), Rome (6), Cape Town (6), and Brisbane (5). The strong representation of Europe and North America reflects their prominence in the data, while the frequent mentions of Australia, South Africa, and New Zealand suggest a focus on biodiversity hotspots or ecological studies. Overall, the dataset spans a diverse range of geographical entities, from continents and regions to countries and cities, emphasizing a global perspective with notable references to the Mediterranean, Asia, and locations such as California and Queensland. Browse the data used for analysis in this paragraph here.

Furthermore, the extracted information provides a comprehensive view of **terrestrial, marine, and aquatic ecosystems**, reflecting their ecological and scientific significance. Terrestrial ecosystems (93) are the most frequently mentioned, with grasslands (42), forests (45), and agricultural landscapes (47) prominently featured, emphasizing biodiversity and land use. Mediterranean ecosystems (37) and tropical ecosystems (26) highlight climate-specific regions, while environments such as riparian (14), alpine (8), and sub-Antarctic (8) systems underscore specialized ecological contexts. Urban ecosystems (46) also feature prominently, reflecting the interplay between natural and human-modified systems. Among ma-

rine ecosystems, the Mediterranean Sea (71) is a standout, alongside coral reefs (8), seagrass meadows (3), and regions like the Baltic Sea (12), showcasing marine diversity. Aquatic ecosystems, particularly freshwater systems (199), are strongly represented, with mentions of wetlands (40), lake ecosystems (59), and riverine ecosystems (36), as well as transitional zones like estuarine ecosystems (35) and coastal wetlands (10), which bridge freshwater and marine environments. Collectively, these ecosystems highlight the dataset's holistic representation of ecological diversity. The full list of the data used for this analysis is here.

Regarding **habitat as a subcomponent of ecosystem** extraction, the following insights were obtained: The dataset reveals nuanced relationships where habitat names and ecosystem types differ across aquatic, marine, and terrestrial environments, showcasing ecological diversity and complexity. In aquatic systems, examples include the *pelagic zone* linked to the lake ecosystem, *ballast water* associated with the marine ecosystem, and *estuarine habitat* tied to the estuarine ecosystem, while unique pairings such as *wetland habitat* with wetland ecosystem and *riverine habitat* with freshwater ecosystem further highlight habitat specificity. In marine environments, *kelp beds* are tied to rocky subtidal ecosystems, and *mussel beds* are associated with the “intertidal zone,” emphasizing distinct ecological contributions. Human-modified habitats, such as *artificial coastal defense structures* linked to *Biogenic reefs*, underscore the role of anthropogenic features, while connections like *maerl habitat* and *seaweed canopies* to coastal ecosystems highlight the contributions of unique marine features. Terrestrial systems also show diverse relationships, including *forest habitat* tied to forest ecosystem, *soybean fields* linked to agricultural ecosystem, and *forest understory* associated with “deciduous forest,” reflecting agricultural and forest-specific ecological roles. Unique pairings, such as *sandy coastal habitats* with “coastal dunes” and *orchards* with orchard ecosystem, emphasize distinct features, while *urban areas* and *domestic gardens* tied to urban ecosystems reflect the ecological importance of human-modified environments. Collectively, this illustrates the dataset's detailed representation of habitat-ecosystem relationships, highlighting their interconnectedness and ecological significance in diverse environments.

The corpus of invasion biology papers reveals a rich diversity of relation types, reflecting the interdisciplinary nature of the field. Predominant categories include **ecological** relations, such as **ecological/anthropogenic**, **ecological/economic**, and **socio-ecological**, which highlight the interactions between natural systems and human influences. Among these, ecological relations are the most prevalent, with **invasion** (814) dominating, followed by **competition** (429), **impact** (349), and **predation** (301), underscoring key processes driving species interactions and environmental changes. Additional relations like **colonization** (179), **distribution** (179), and **habitat preference** (123) emphasize the dynamics of species spread, habitat utilization, and behavioral patterns. Complementing this, **biological relations** such as **predation** (177), **parasitism** (151), and **hybridization** (74) focus on specific ecological interactions, while associations like **pollination** (25) and **symbiosis** (25) highlight cooperative and movement-based dynamics. Furthermore, **physical relations**, including **location**, **transport**, and **introduction location**, emphasize spatial and movement dynamics critical to understanding alien species impacts and habitat transformations. Lastly, **anthropogenic relations** are dominated by **introduction** (157), along with **introduction pathway** (45) and **transportation** (14), underscoring the significant role of human activities in species dispersal and ecological consequences. Collectively, these relations capture the multifaceted processes and interactions that define invasion biology.

The wealth of information mined from a fully unsupervised IE task conducted by an LLM highlights the immense potential of large language models as tools for researchers and practitioners in ecological sciences and beyond. These models have the potential to serve as valuable assistants in tasks such as systematic or scoping reviews. The results presented here represent only a fraction of the insights that can be derived from the extracted information in our large corpus of over 10,000 papers. To support further research, we have released the corpus at <https://doi.org/10.5281/zenodo.13956882>.

This work aligns with the paradigm of open information extraction (OIE) (Etzioni et al., 2008; Fader et al., 2011; Etzioni et al., 2011), which traditionally relies on syntactic patterns in text, such as verbs for relations and (subject, object) depen-

dency structures. However, the advanced semantic comprehension capabilities of LLMs have far surpassed traditional OIE methods. By extending beyond syntax into the realm of semantics, LLMs provide highly practical solutions for analyzing complex relationships in large-scale corpora.

#### 4 Recommendations for Future Work

The semantic web community presents technologies and concepts that may supply direction to the future work on this topic. Ontologies have been a focal point of research in the domain for decades, however, the expertise required for their use has presented a barrier to outside establishment. With the rise in prominence of LLMs in both research and application, a new branch of research has emerged with the goal to investigate the interplay of ontologies and LLMs for more precise IE and to facilitate creating linked data. This branch currently covers the following topics: How can LLMs support ontology and knowledge graph construction? (Komminen et al., 2024) How can LLMs for question answering be improved through ontology support? (Allemang and Sequeda, 2024) How can LLMs support ontology learning from texts? (Babaei Giglou et al., 2023, 2024) How can representation learning of LLMs be enhanced by ontologies? (Ronzano and Nanavati, 2024)

The concept of ontologies, defined as a *formal, explicit specification of a shared conceptualisation* (Studer et al., 1998), highlights their role in enabling a *shared understanding of a domain* and reflecting *consensus about domain knowledge*. This study builds on these principles, exploring how information can be extracted from unstructured texts and structured using semantic models. Future work can extend this approach by integrating ontological knowledge through schema-driven IE, addressing key questions: What information should be provided to LLMs to improve extraction, and where in the workflow is it most effective? How can LLMs be precisely guided (Cauffman et al., 2024) toward achieving semantic modeling objectives? Finally, how does shared human consensus in a domain align with the consensus extracted by LLMs trained on human knowledge? Exploring these questions in future work can bridge the gap between human and machine understanding of domain knowledge.

We identify several applications of ontologies in IE workflows to enhance LLM performance.

Incorporating domain-specific ontologies during training could improve understanding of domain terms by supplying definitions, properties, and hierarchical relations, aligning the model's comprehension with that of domain experts. Ontologies can also guide semantic modeling by constraining or informing LLMs, improving interoperability and making outputs more actionable. For entity recognition (ER/NER) and relation extraction (RE), pre-informing models about common domain-specific terms and relationships may enhance accuracy by effectively teaching the model "what to look for." Additionally, retrieval-augmented generation (RAG) integrates knowledge bases, such as ontologies or knowledge graphs, into QA systems, reducing hallucinations and improving performance, as demonstrated in domains like biomedicine (Soman et al., 2024).

Finally, these advantages need to be balanced with better understanding of what such constraints on LLMs might mean for their use in domains which have less well-structured general knowledge, perhaps because they are still emerging, or perhaps because the knowledge in those domains is changing rapidly. In such cases, there will be a need for further work to understand whether and to what extent using ontologies to constrain LLMs may also constrain our ability to track emerging or rapidly changing knowledge.

#### 5 Conclusion

This study demonstrates the potential of LLMs for advancing IE in invasion biology by mining key entities such as species, locations, habitats, and ecosystems from scientific literature. By developing and applying a standardized semantic schema, we showcased the ability of LLMs to structure complex ecological data, providing a foundation for enhancing workflows in ecological research. Using a two-stage approach, the specialize stage extracts detailed, context-specific structures, while the generalize stage integrates these into a flexible schema balancing specificity and generality. This method addresses the complexity of ecological systems, enabling structured representation of nuanced information. The release of the dataset and schema enable refining extraction methods, integrating ontologies, and exploring broader ecological applications. This research underscores the utility of LLMs as tools for bridging unstructured data and structured knowledge in ecology.



## Acknowledgments

The experiments reported in this paper using the OpenAI models were supported by the SCINEXT project grant (BMBF, German Federal Ministry of Education and Research, Grant ID: 01IS22070) and by the TIB Leibniz Information Centre for Science and Technology, Hannover, Germany. This work is also conducted in the context of the ZiF interdisciplinary research group “Mapping Evidence to Theory in Ecology,” hosted at Bielefeld University.

## Limitations

While this study highlights the potential of LLMs for IE in invasion biology, certain limitations remain. The extracted entities and relations were not evaluated against a gold-standard dataset, making it difficult to quantify precision and recall. A future inter-annotator agreement (IAA) study on a subset of the corpus (e.g., 20%) or a qualitative error analysis could enhance its reliability for researchers. Our approach also relies solely on OpenAI GPT-4o, without comparing alternative LLMs or prompting strategies, such as chain-of-thought prompting, which may improve extraction accuracy. Additionally, potential data contamination (?) remains a concern, as LLMs may reproduce information seen during pre-training rather than extracting it anew. A systematic comparison against pre-training corpora would help assess this effect.

## References

- Stein Aerts, Diether Lambrechts, Sunit Maity, Peter Van Loo, Bert Coessens, Frederik De Smet, Leon-Charles Tranchevent, Bart De Moor, Peter Marynen, Bassem Hassan, et al. 2006. Gene prioritization through genomic data fusion. *Nature biotechnology*, 24(5):537–544.
- Dean Allemang and Juan Sequeda. 2024. Increasing the llm accuracy for question answering: Ontologies to the rescue! *arXiv preprint arXiv:2405.11706*.
- Xavier Amatriain, Ananth Sankar, Jie Bing, Praveen Kumar Bodigutla, Timothy J Hazen, and Michael Kazi. 2023. Transformer models: an introduction and catalog. *arXiv preprint arXiv:2302.07730*.
- Isabelle Augenstein, Mrinal Das, Sebastian Riedel, Lakshmi Vikraman, and Andrew McCallum. 2017. Semeval 2017 task 10: Scienceie-extracting keyphrases and relations from scientific publications. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 546–555.
- Mohamed-Amine Baazizi, Clément Berti, Dario Colazzo, Giorgio Ghelli, and Carlo Sartiani. 2020. Human-in-the-loop schema inference for massive json datasets. In *EDBT 2020-23rd International Conference on Extending Database Technology*, pages 635–638. OpenProceedings.org.
- Mohamed-Amine Baazizi, Housseem Ben Lahmar, Dario Colazzo, Giorgio Ghelli, and Carlo Sartiani. 2017. Schema inference for massive json datasets. In *Extending Database Technology (EDBT)*.
- Hamed Babaei Giglou, Jennifer D’Souza, and Sören Auer. 2023. Llm4ol: Large language models for ontology learning. In *International Semantic Web Conference*, pages 408–427. Springer.
- Hamed Babaei Giglou, Jennifer D’Souza, and Sören Auer. 2024. Preface for llms4ol 2024: The 1st large language models for ontology learning challenge at the 23rd iswc. *Open Conference Proceedings*, 4:1–2.
- Arthur Brack, Jennifer D’Souza, Anett Hoppe, Sören Auer, and Ralph Ewerth. 2020. Domain-independent extraction of scientific concepts from research articles. In *European Conference on Information Retrieval*, pages 251–266. Springer.
- Marc Brinner, Tina Heger, and Sina Zarriess. 2022. Linking a hypothesis network from the domain of invasion biology to a corpus of scientific abstracts: The INAS dataset. In *Proceedings of the first Workshop on Information Extraction from Scientific Publications*, pages 32–42, Online. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Phillip Cassey, Pablo García-Díaz, Julie L Lockwood, and Tim M Blackburn. 2018. Invasion biology: searching for predictions and prevention, and avoiding lost causes. In *Invasion biology: hypotheses and evidence*, pages 3–13. CAB International Wallingford UK.
- J Harry Caufield, Harshad Hegde, Vincent Emonet, Nomi L Harris, Marcin P Joachimiak, Nicolas Matentzoglou, HyeonSik Kim, Sierra Moxon, Justin T

- Reese, Melissa A Haendel, et al. 2024. Structured prompt interrogation and recursive extraction of semantics (spires): A method for populating knowledge bases using zero-shot learning. *Bioinformatics*, 40(3):btae104.
- Danilo Dessì, Francesco Osborne, Diego Reforgiato Recupero, Davide Buscaldi, Enrico Motta, and Harald Sack. 2020. Ai-kg: an automatically generated knowledge graph of artificial intelligence. In *The Semantic Web—ISWC 2020: 19th International Semantic Web Conference, Athens, Greece, November 2–6, 2020, Proceedings, Part II 19*, pages 127–143. Springer.
- Jennifer D’Souza, Sören Auer, and Ted Pedersen. 2021. SemEval-2021 task 11: NLPContributionGraph - structuring scholarly NLP contributions for a research knowledge graph. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 364–376, Online. Association for Computational Linguistics.
- Jennifer D’Souza. 2023. A catalog of transformer models.
- Jennifer D’Souza. 2024. Agriculture named entity recognition—towards fair, reusable scholarly contributions in agriculture. *Knowledge*, 4(1):1–26.
- Jennifer D’Souza and Sören Auer. 2022. Computer science named entity recognition in the open research knowledge graph. In *International Conference on Asian Digital Libraries*, pages 35–45. Springer.
- Jennifer D’Souza, Salomon Kabongo, Hamed Babaei Giglou, and Yue Zhang. 2024. Overview of the clef 2024 simpletext task 4: Sota? tracking the state-of-the-art in scholarly publications. In *Working Notes of the Conference and Labs of the Evaluation Forum (CLEF 2024)*, pages 3163–3173.
- Oren Etzioni, Michele Banko, Stephen Soderland, and Daniel S Weld. 2008. Open information extraction from the web. *Communications of the ACM*, 51(12):68–74.
- Oren Etzioni, Anthony Fader, Janara Christensen, Stephen Soderland, et al. 2011. Open information extraction: The second generation. In *Twenty-Second International Joint Conference on Artificial Intelligence*. Citeseer.
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. Identifying relations for open information extraction. In *Proceedings of the 2011 conference on empirical methods in natural language processing*, pages 1535–1545.
- Kata Gábor, Davide Buscaldi, Anne-Kathrin Schumann, Behrang QasemiZadeh, Haifa Zargayouna, and Thierry Charnois. 2018. SemEval-2018 task 7: Semantic relation extraction and classification in scientific papers. In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 679–688, New Orleans, Louisiana. Association for Computational Linguistics.
- John Giorgi, Xindi Wang, Nicola Sahar, Won Young Shin, Gary D Bader, and Bo Wang. 2019. End-to-end named entity recognition and relation extraction using pre-trained language models. *arXiv preprint arXiv:1912.13415*.
- Jonathan M Jeschke and Tina Heger. 2018. *Invasion biology: hypotheses and evidence*. CAB International.
- Salomon Kabongo, Jennifer D’Souza, and Sören Auer. 2021. Automated mining of leaderboards for empirical ai research. In *International Conference on Asian Digital Libraries*, pages 453–470.
- Petr Knoth, Drahomira Herrmannova, Matteo Cancellieri, Lucas Anastasiou, Nancy Pontika, Samuel Pearce, Bikash Gyawali, and David Pride. 2023. Core: a global aggregation service for open access papers. *Scientific Data*, 10(1):366.
- Vamsi Krishna Kommineni, Birgitta König-Ries, and Sheeba Samuel. 2024. From human experts to machines: An llm supported approach to ontology and knowledge graph construction. *arXiv preprint arXiv:2403.08345*.
- Gang Li, Karen E Ross, Cecilia N Arighi, Yifan Peng, Cathy H Wu, and K Vijay-Shanker. 2015. mirtex: a text mining system for mirna-gene relation extraction. *PLoS computational biology*, 11(9):e1004391.
- Haoyang Liu, M Janina Sarol, and Halil Kilicoglu. 2021. Uiuc\_bionlp at semeval-2021 task 11: A cascade of neural models for structuring scholarly nlp contributions. *arXiv preprint arXiv:2105.05435*.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232.
- Daniel Mietchen, Jonathan M. Jeschke, Maud Bernard-Verdier, Tina Heger, Camille Musseau, and Steph Tyszka. 2024. Invasion biology corpus 2024-07.
- Riccardo Miotto, Fei Wang, Shuang Wang, Xiaoqian Jiang, and Joel T Dudley. 2018. Deep learning for healthcare: review, opportunities and challenges. *Briefings in bioinformatics*, 19(6):1236–1246.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Francesco Ronzano and Jay Nanavati. 2024. Towards ontology-enhanced representation learning for large language models. *arXiv preprint arXiv:2405.20527*.
- Helen E Roy, Anfal Pauchard, Peter Stoett, Tanara Renard Truong, Sven Bacher, Bella S Galil, Philip E Hulme, Tohru Ikeda, Kavileveetil Sankaran,

- Melodie A McGeoch, et al. 2023. Ipbes invasive alien species assessment: summary for policymakers. *IPBES*.
- Mahsa Shamsabadi, Jennifer D'Souza, and Sören Auer. 2024. Large language models for scientific information extraction: An empirical study for virology. In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 374–392, St. Julian's, Malta. Association for Computational Linguistics.
- Karthik Soman, Peter W Rose, John H Morris, Rabia E Akbas, Brett Smith, Braian Peetoom, Catalina Villouta-Reyes, Gabriel Ceron, Yongmei Shi, Angela Rizk-Jackson, et al. 2024. Biomedical knowledge graph-optimized prompt generation for large language models. *Bioinformatics*, 40(9):btae560.
- Rudi Studer, V Richard Benjamins, and Dieter Fensel. 1998. Knowledge engineering: Principles and methods. *Data & knowledge engineering*, 25(1-2):161–197.
- Zhong-Yi Wang and Hong-Yu Zhang. 2013. Rational drug repositioning by medical genetics. *Nature biotechnology*, 31(12):1080–1082.
- XueZhong Zhou, Jörg Menche, Albert-László Barabási, and Amitabh Sharma. 2014. Human symptoms–disease network. *Nature communications*, 5(1):4212.