BEAR: Value-First Ontology Engineering Framework for Business Ecosystem Analysis and Representation

Alican Tüzün[1,2,*], Nick Bassiliades[1], Herbert Jodlbauer[2] and Georgios Meditskos[1]

[1]School of Informatics, Aristotle University of Thessaloniki, Thessaloniki, Greece
[2]Josef Ressel Centre for Data-Driven Business Model Innovation, University of Applied Sciences Upper Austria, Wehrgrabengasse 1-4, 4400, Steyr, Austria

Abstract

High-quality ontology engineering traditionally prioritizes complete, reusable domain models. While effective for broad reuse, this “ontology-first” approach can misalign with the needs of strategic decision makers, who need targeted, actionable insights on constrained timelines. This paper introduces a value-first framework that inverts this process, beginning with the strategic goals, jobs, and knowledge gaps of business leaders to generate lean, purpose-built knowledge graph that delivers immediate value. In a pilot project with CompanyA, we applied this framework to the wind energy ecosystem, successfully answering 15 distinct knowledge questions. To demonstrate this, we focus on one such question out of 15, analyzing data from 35 companies collected at WindEnergy Hamburg 2024. Our findings show that this approach not only answers knowledge questions effectively through tailored visualizations but also uncovers critical blind spots—such as the intermediary roles of consulting firms—that conventional business ecosystem analyses would necessarily miss.

Keywords

Ontology Engineering, Business Ecosystem Analysis, Knowledge Graph Engineering, Value-First Ontology Engineering, Strategic Decision Support

1. Introduction

A business ecosystem is not just a collection of isolated entities but a dynamic system, much like a biological one [1, 2]. Its dynamism emerges from its active organizations (e.g., corporations, non-profits) interconnected by shared goals, value propositions, and relationships, creating a causal unity. This unity, however, presents a double-edged sword. On the one hand, it can foster innovation and shared success for these actors; on the other, disruptions [3] within the ecosystem (e.g., a key player’s failure or an innovation) can significantly affect the entire ecosystem [1, 2].

Therefore, overlooking relationships or the implicit roles of certain actors within these ecosystems [1]—what we term “blind spots”— during the decision-making process, means failing to navigate the dangers inherent in this double-edged sword. This danger is not merely an academic oversight; such blind spots can directly hinder practical strategic activities, obscure market opportunities, and leave an organization vulnerable in its competitive position, which could make or break its future [1, 2]. Therefore, effectively uncovering such blind spots demands practical, structured approaches, where Ontology and Knowledge Graph Engineering (OKGE) offer significant promise [4, 5, 6].

Current OKGE methodologies, however, create a fundamental mismatch with the needs of strategic decision-makers (e.g., Business Development Managers, Chief Innovation Officers). For example,

16th Workshop on Ontology Design and Patterns (WOP 2025), 20th International Semantic Web Conference (ISWC 2025), November 2-6, 2025, Nara, Japan

*Corresponding author.
alican.tuezuen@fh-steyr.at (A. Tüzün); nbassili@csd.auth.gr (N. Bassiliades); herbert.jodlbauer@fh-ooe.at (H. Jodlbauer); gmeditsk@csd.auth.gr (G. Meditskos)
https://github.com/T-Z-N (A. Tüzün); https://csd.auth.gr/en/member/nick-bassiliades (N. Bassiliades); https://www.fh-ooe.at/en/about-us/people/herbert-jodlbauer (H. Jodlbauer); https://csd.auth.gr/en/member/georgios-meditskos (G. Meditskos)
0009-0009-8017-5487 (A. Tüzün); 0000-0001-6035-1038 (N. Bassiliades); 0000-0001-9818-7931 (H. Jodlbauer); 0000-0003-4242-5245 (G. Meditskos)

established OKGE frameworks (e.g., METHONTOLOGY [7]) offer valuable methodologies for engineering comprehensive, reusable domain models. To serve broad communities, their process is logically organized around domain completeness and user needs. Nevertheless, strategic decision-makers face a different challenge entirely. They need targeted competitive intelligence that directly answers specific business questions (e.g., *Can we offer additional services to our potential or current customers?*) to make the decision quickly, rather than comprehensive domain representations [8, 9].

This mismatch represents a fundamental quality and expectation gap in how OKGE creates value. Current methodologies treat strategic insight as an afterthought—an emergent property of a technically sound model. The philosophy is: build a complete, consistent ontology first, assuming strategic value will eventually follow. Strategic decision-makers need the inverse: a targeted ontology, designed specifically around their knowledge gaps, to help them get their job done [3], whether that job is driving revenue growth, fostering innovation, and gaining a competitive advantage [3, 10].

Therefore, bridging this gap requires a paradigm shift. Ontology engineering must explicitly organize the development process around strategic business objectives as the primary design driver, rather than treating strategic value as an emergent consideration. This demands new frameworks that explicitly translate strategic goals directly into ontology design decisions— from initial value definition to final value delivery—to get decision makers' job done [3].

This pressing pragmatic need leads to our research question: *How can strategic business objectives be systematically operationalized as the foundational organizing principle for ontology and knowledge graph engineering to enable competitive intelligence in dynamic business ecosystems?*

To address this, we introduce BEAR [11], a value-first framework that makes three key contributions. First, it establishes strategic business objectives as the foundational organizing principle for OKGE, putting it in a value-driven context. Second, it provides a systematic process to translate those objectives into ontological design decisions. Moreover, third, it reveals critical blind spots that traditional business ecosystem analysis in the literature misses. In the following sections, we will discuss related work, outline BEAR's core principles, and demonstrate its value through a case study in the wind energy ecosystem, key findings, and future directions.

2. Background and Related Work

As we have argued in the introduction, analyzing business ecosystems requires bridging two distinct research domains: high-level business strategy with ontology and knowledge graph engineering (OKGE). This section builds the arguments for our approach in two steps. First, we review the state of business ecosystem literature, to identify a critical gap: the lack of methods that are both semantically rich and empirically grounded. Second, we argue that while OKGE is perfectly suited to fill this gap, current ontology-centric methodologies are not explicitly designed for strategic analysis. Therefore, the following review will show a consistent pattern of treating strategic insight as a secondary, emergent property rather than the primary design driver, thereby establishing the critical knowledge gap that the BEAR framework is designed to fill.

2.1. Approaches to Business Ecosystem Analysis

To analyze business ecosystems, recent data-driven methodologies utilize natural language processing and text mining on large unstructured text corpora (e.g., company reports) [12]. These methods typically identify relevant entities and construct interactive network visualizations based on the entity co-occurrence within the source documents. However, relying on textual co-occurrence has significant semantic limitations. For example, explicit intensions regarding the relationships between entities are often missed [13] or, in the best case, inferred statistically from textual proximity (e.g, cosine similarity) [12, 14, 15]. Therefore, this paradigm is mostly statistical/syntactic, lacking intensions crucial for semantic analysis.

Beyond data-driven approaches, structured conceptual modeling offers alternative ways to analyze business ecosystems, focusing on intensions [1, 16, 17, 18]. Methodologies like e3 value provide foundations for modeling economic value exchanges, enabling the analysis of financial viability and value flows within defined networks [18]. Some authors further explored the usability of such conceptual modeling, for example, through tangible interfaces that aim to map complex modeling languages and practitioner needs [19]. While these approaches provide valuable frameworks for understanding value networks, they heavily focus on intensions rather than instances and their extensions [1, 16, 17, 18, 19], which is necessary for deductive inference for deeper ecosystem analysis.

2.2. Ontology Engineering Methodologies: From Domain-Centric to Value-Centric

Established OKGE frameworks, such as METHONTOLOGY, offer robust methods for domain analysis but are fundamentally product-centric [7]. They focused on defining the ontology, treating strategic value as an indirect outcome rather than a primary engineering driver. This paradigm is evident in their use of ontology requirements specification documents, which detail the ontology's intended users, and Competency Questions (CQs), much like a blueprint for a software product [7, 20]

This product-centric focus persists even in collaborative, efficiency-driven methodologies with disintermediating efforts of ontology engineers [21]. UPON lite, which is a light version of the UPON methodology [22], rightly critiques traditional ontology building (e.g., [7]) as too expensive and time-consuming. Even though we agree on this criticism, the method's focus is still on the ontology itself, evident from its first step: defining the domain terminology [21]. Similarly, frameworks like LOT [23] aim for greater precision in requirements by adding more granular details and options, but the goal remains the same: to specify what the ontology must represent. A business leader, however, is still left asking, "So what?".

The Extreme Design (XD) methodology is notable for explicitly incorporating business value into its design philosophy, inspired by Extreme Programming practices [24]. Its extreme lightweight ontology, maintenance, and prototyping principles aim to deliver quick business value, rather than building for the abstract future [24]. XD tries to capture this value through planning, where customers define their needs via desired features and CQs. However, this is where the methodology's business value philosophy disconnects from its product-centric process. While customers are asked to define "business value", the mechanism for this is still the CQ, a tool for specifying the product's features (See Table 1). Furthermore, XD does not explicitly describe how to derive the initial "baseline ontology" [25] from strategic goals. Ultimately, despite its aims and starting point, the process later focuses on building the ontology itself, rather than building the ontology around the given strategic goals and knowledge gaps of the stakeholders.

Even highly innovative structural approaches, such as the Modular Ontology Modelling methodology (MOMo), illustrate this focus on the ontology as the end product. Building upon the Extreme Design Methodology, MOMo's guiding principle is not the traditional taxonomical (is-a) hierarchy. Instead, it prioritizes the modularity of the ontology, viewing each module as a part of the whole [26]. Despite this novel and innovative approach to ontology engineering, the first two steps, as well as the example use case descriptions of the MOMo workflow, indicate that the methodology still focuses on the ontology itself, seeing it as the end product [26]. This journey through the current landscape, from traditional to recent, reveals a clear and consistent theme. If we generalize, they ask the question, "What must this ontology be able to represent?" This inevitably forces the strategic decision maker to adapt to the model's structure.

Therefore, a critical gap remains for a methodology that inverts this process. We argue for a shift away from defining a model's features and toward delivering its value. In this context, we define value not as a model's technical completeness or reusability but as the degree to which it closes a specific, strategic knowledge gap for a decision-maker to get their job done and reach their or their organization's business goals. Therefore, an ontology became a means to an end rather than an end in itself. These arguments lead us to ask different type of question: "What strategic question must this ontology answer to get the job of decision maker done?"This shift is the foundation of the BEAR framework, which we introduce next.

Table 1
The Paradigm Shift from Competency Questions (CQs) to Knowledge Questions (KQs). This table contrasts the traditional, product-centric view of ontology requirements (CQs) with BEAR's value-first approach (KQs), highlighting fundamental differences in purpose, audience, and strategic outcome.

Criterion	Competency Question (CQ) *(The Traditional, Ontology-First View)*	Knowledge Question (KQ) *(The BEAR, Value-First View)*
Primary Purpose	To **verify the functional scope** of the ontology. It serves as a technical "spec sheet" for the model.	To **articulate a strategic knowledge gap** of the business. It serves as the "problem statement" for the stakeholder.
Primary Audience	**Ontology Engineers & Knowledge Modelers.** It guides their implementation and testing.	**Business Decision-Makers & Strategists.** It frames their problem and sets the success criteria for the project.
Point of Origin	Derived from domain analysis and use cases. Answers: *"What must our ontology be able to represent?"*	Derived from high-level business goals and "Jobs-To-Be-Done" [3]. Answers: *"What must we know to succeed?"*
Object of Inquiry	The **ontology itself**. The question is about the model's capabilities (its classes, properties, axioms).	The **business ecosystem**. The question is about real-world dynamics, relationships, and emergent structures.
Nature of Question	**Factual & Verificational.** Typically asks "What" or "List". It can often be answered with a direct SPARQL query.	**Analytical & Exploratory.** Typically asks "How" or "Why". The answer is an insight revealed through reasoning and visualization.
Role in Methodology	A **requirement** for the design phase. It helps define the necessary components of the ontology.	The **foundational driver** for the entire process, dictating data collection, ontology design, and the final value delivery.
Resulting Ontology	Tends to produce a **comprehensive, reusable domain ontology**. The ideal is logical completeness and correctness.	Produces a **lean, purpose-built "Seed Ontology"**. The ideal is pragmatic sufficiency to answer the KQ.
Associated Paradigm	**Product-Centric.** The focus is on building a robust knowledge *product*. Strategic value is an expected *emergent property*.	**Value-First.** The focus is on delivering a specific strategic *answer*. The ontology is a *means to that end*.
Core Trade-off	Risk of high development cost and creating a technically perfect model that is **strategically irrelevant** (the "so what?" problem).	Risk of creating **siloed, non-reusable graphs** for each KQ, sacrificing long-term asset building for short-term answers.
Example	"Can the ontology represent a product delivery relationship between two companies?" *or* "Can we list all the bearing manufacturers that deliver to generator manufacturers?"	"How do specific companies establish their positions through product/service delivery interactions within the wind energy ecosystem?"

3. The BEAR Framework

The BEAR framework has three core processes: Value Definition, Creation, and Delivery. This section describes the principles and workflow of each process in a sequence. The next section presents a detailed case study to provide a clear illustration of this framework in action.

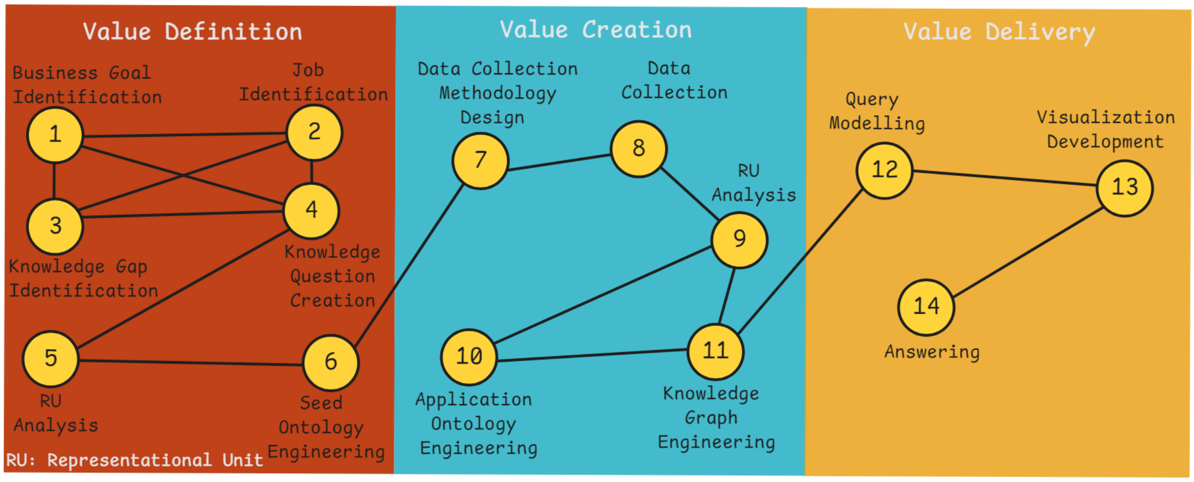

Figure 1: The BEAR framework translates business goals into strategic value delivery. Undirected edges in the figure represent the iterative nature of the subprocesses and temporal proximity. The process moves from left to right through three phases. (1) The value definition process begins by identifying a stakeholder's business goals, jobs, and knowledge gaps to formulate a specific knowledge question and formalize it into a seed ontology. (2) Value Creation process uses that seed ontology to guide data collection and engineer a targeted knowledge graph. (3) Value Delivery process delivers the answer to a knowledge question through query modeling and tailored visualizations, revealing strategic insights.

3.1. Value Definition: From Business Goal to Seed Ontology

BEAR's value definition process begins by treating decision makers as customers. The first process is to identify their high-level business goal—goals that are strategic, informal objectives they express in natural language. From these goals, BEAR defines the stakeholder's specific "Jobs To Be Done (JTBD)"—a progress a customer is trying to make in a given circumstance to advance their goals [3, 27]. By identifying the job the decision-maker is trying to get done (e.g., "assess a new market opportunity"), BEAR can then identify the precise knowledge gap that prevents them from completing that job successfully, facilitating progress toward their goals.

Through collaborative elicitation (e.g., workshops, interviews, meetings), we find these business goals, jobs, and knowledge gaps, subsequently formulating the gaps into single, consensus-driven knowledge questions (KQ). Following iteratively refining the KQ with decision makers, BEAR engineers analyze it to identify the core Representational Units (RUs), which are the minimal semantic components (classes, properties, relationships), in our context, needed to answer the KQ [28]. These RUs are then formalized into a lean Seed Ontology (SO) similar to baseline ontology [24, 25] in OWL [29], leveraging its expressive power and formal semantics.

This SO is neither a comprehensive domain nor an application ontology. Instead, it formally represents what the stakeholder needs to know to get their job done. It functions as a lightweight, reusable pattern for value-first modeling, emphasizing pragmatic sufficiency over exhaustive domain coverage, much like ontology design patterns (ODPs) promote reuse and modularity in ontology engineering [26]. In a business context, this SO provides a reusable blueprint that ontology engineers can "hire" to address similar knowledge gaps across different ecosystems, ensuring alignment with stakeholder jobs-to-be-done [27].

3.2. Value Creation: From Seed Ontology to Knowledge Graph

In the value creation process, SO becomes the template for designing the data collection methodology. It ensures we only target the evidence that addresses the stakeholder knowledge gap. Based on the SO and stakeholder agreement, we specify the exact data sources (e.g., expos, websites, databases, social media) and collection protocols (e.g., interviews, surveys, web scraping) for the value creation process.

Once we collect the data, BEAR analyzes RUs within the collected evidence and iteratively maps them to the developing SO. During this mapping process, BEAR employs a form of reification [30] to handle evidence where specific information about an entity is ambiguous or incomplete (e.g., an entity is mentioned by company type but not by company name).

This iterative mapping evolves the lean SO into a more comprehensive application ontology, and we argue, potentially towards a reference ontology through continued iterations. Crucially, deductive reasoning, employing standard OWL reasoners [29], plays a significant role in consistency checking through this SO evolution.

3.3. Value Delivery: From Graph to Answering the Knowledge Question

BEAR's third and final process, value delivery, answers the stakeholder's KQ. Following the business ecosystem literature, we develop interactive, tailored visualizations [12, 14]. BEAR enables these visualizations and allows access to the enriched knowledge by modeling SPARQL queries. These queries leverage reasoners to retrieve explicitly asserted and implicitly inferred information within the knowledge graph.

Importantly, these inferences are not mere technical artifacts. They are the engine for revealing strategic insights that decision makers are looking for and exposing critical blind spots— the core value BEAR was designed to deliver.

To deliver these insights effectively, BEAR employs and advocates custom, flexible visualizations built with libraries like D3.js [31]. We argue that standard off-the-shelf ontology visualization tools are too rigid to answer specific knowledge questions, as they are often designed to display the overall schema, rather than enable exploration of instance-class level interactions [32, 33]. In contrast, a tailored visualization can provide more than just a picture of the data; it is a critical mechanism for delivering value— an answer engineered to close the knowledge gap to get the stakeholder's job done to reach business goals.

4. Case Study: Applying BEAR to the Wind Energy Ecosystem

To validate and illustrate the BEAR framework, we applied it in a pilot project with CompanyA to a real-world strategic challenge within the wind energy business ecosystem. The project followed the three core processes of the BEAR framework.

Because we advocate for open science and reproducibility, we have made the research artifacts from this pilot publicly available in our GitHub repository [11]. The repository includes the semi-structured survey, the resulting knowledge graph, its visualization, and the SPARQL queries used for the selected knowledge question. Proprietary materials, such as raw stakeholder discussions and the visualization source code, have been excluded.

4.1. Applying Value Definition Process in the Wind Energy Ecosystem Analysis

The value definition process began with a collaborative workshop with senior leaders at CompanyA. Their high-level business goal was clear: to drive revenue growth in the increasingly crowded wind energy ecosystem. Nevertheless, they were struggling to reach that goal. Traditional market analysis and business development efforts felt like they were competing against the luck [27]—trying to innovate their business models, often with no real way to predict their organization's success.

At this moment of struggle, we helped them articulate the specific jobs to be done, which prevented them from making progress as advocated in the jobs to be done theory [27]. The real job they needed to get done was: "Help us see the hidden relationships of our ecosystem so we can confidently identify new, high-margin service opportunities and stop wasting resources and company capabilities on low-probability bets." This job was not the only one we identified; however, for this paper, we focus on this one. Therefore, within this context, getting this job done was one of the ways to achieve their broader business goal of revenue growth.

Beginning to get this job done, we had to move from their abstract struggle to concrete knowledge gaps. Here we argue that, to get this job done, they also might need human resources, special equipment, and many other resources; however, another important point relevant to the ontology engineering is the knowledge gap they had: "company positions within wind energy ecosystem achieved through product/service delivery interactions".

As competency questions bridges the gap between the ontology engineers and stakeholders [7, 23, 34], we formalized the knowledge gap into a Knowledge Question (KQ), however acknowleding the difference (See Table 1), which would serve as the foundation for the rest of the BEAR process: "How do specific companies establish their positions through product/service delivery interactions within the wind energy ecosystem?".

Analyzing this KQ revealed related Representational Units (RUs), such as "Company", "Product Delivery Interaction", and "Service Delivery Interaction". A systematic analysis was used to distill the KQ into its core semantic components (Table 2). We then formalized these RUs in OWL2 to engineer the Seed Ontology (SO), with design decisions guided explicitly by the Value-First Principle (Table 3).

Table 2
Systematic RU Analysis Process Applied to Wind Energy KQ

Step 1: Left-to-Right Parsing	
Input:	Complete KQ
Output:	17 candidate words
Process:	Sequential word extraction: *How, do, specific, companies, establish, their, positions, through, product, service, delivery, interactions, within, the, wind, energy, ecosystem*
Step 2: POS Filtering	
Input:	17 candidate words
Output:	7 selected terms
Criteria:	Singular nouns and noun phrases representing domain entities, or relationships: *companies, positions, product, service, product delivery interaction, service delivery interaction, wind energy ecosystem*
Step 3: RU Selection & Normalization	
Input:	7 selected terms
Output:	3 RU decisions
Principle:	Model only concepts essential for answering the KQ: *Company, product delivery interaction, service delivery interaction*

4.2. Applying Value Creation Process: From Seed Ontology to Application Ontology

For this application, we used the SO to guide the data collection methodology creation process directly. We designed a semi-structured survey for rapid, open-ended data acquisition at industrial expos [11]. With stakeholder approval, we utilized this survey at WindEnergy Hamburg 2024, which is one of the largest wind energy expos in the world [35]. This data collection process in the expo yielded 37 filled surveys from 35 companies. After anonymizing the data, we iteratively mapped the collected RUs from these data back onto the SO.

As anticipated, this iterative modeling process uncovered new, more abstract classes not present in our initial SO. For example, we created a wbeo:Operator parent class to unify wbeo:GridOperator and wbeo:WindTurbineOperator (See Figure 2). To manage these emergent abstractions with ontological rigor, we applied the established principle of single inheritance [28]. This process also forced us to handle incomplete data, for which we used reification—creating typed blank nodes for relationship modelling (Figure 3).

Consequently, through this iterative process of mapping and refinement, the SO evolved into the Wind Business Ecosystem Ontology (WBEO) [11]—an application ontology that offers a clear pathway

Table 3
Engineering the Seed Ontology with the Value-First Principle. This table shows the design decisions for modeling key concepts from the Knowledge Question, prioritizing a lean, value-focused ontology over a comprehensive one.

Selected Term (from KQ)	Operationalization in the Seed Ontology	Rationale (BEAR Value-First Principle)
`company`	Modeled as a primary class: `:Company`, with specific types (e.g., ':BearingManufacturer') added as subclasses during modeling.	The company is the central actor. Typing the companies is essential for inferring the nature of their interactions.
`delivery interaction`	Modeled as two core object properties: `:deliversTo` and its inverse, `:receivesFrom`.	This directly captures the primary dynamic described in the KQ. Modeling the action as a relationship between companies is the most efficient way to build the network.
`product`, `service`	**Not modeled as explicit classes**. The specific item being delivered is semantically implied by the defined types of the interacting ':Company' instances.	The strategic goal is to map the value network, not to build a product catalog. Inferring the product from the actor types keeps the SO extremely lean and focused on relational intelligence.
`position`	**Not modeled in the SO**. Treated as an analytical outcome.	A company's position is an emergent property of its network of interactions. It is the *answer* we seek from the final KG, not a concept to model at the start.
`wind energy ecosystem`	**Not modeled as a class in the SO**. It is represented by the entire instance graph of all companies and their interactions.	The ecosystem is the holistic structure of all companies and their interactions. The goal is to represent this emergent structure, not define it as a single entity.

towards a reference ontology for the wind energy domain. Although BEAR advocates the principle of reuse, we developed this SO primarily from RUs due to pilot project constraints and quality concerns of existing domain ontologies [36].

4.3. Applying Value Delivery Process: From Application Ontology to Knowledge Question Answering

In our wind energy application, we first modeled two SPARQL queries [11] to answer KQ. We executed these queries in GraphDB, using its OWL 2 RL (also valid in DL) reasoner [33]. This allowed us to extract both asserted and inferred facts, such as deduced delivery links (See Figure 4 and Figure 3).

Finally, we exported these results as a JSON file and fed them into an interactive visualization developed with D3 [31]. This final tool did more than just displaying the inferred data, it answered the stakeholder's KQ directly, revealing strategic blind spots through features like filtering and granularity adjustments (See Figure 4).

In conclusion, the project resulted in several meetings of the interactive visualization with key decision-makers within different departments at CompanyA. During the meetings, stakeholders could explore the interactive visualization and ask follow-up questions, which led to further insights and discussions about their business ecosystem. The value was directly apparent, leading to two key outcomes: a richer, shared understanding of their business ecosystem and the formal approval of a new pilot project. This second project will test the BEAR framework's utility in a new business context, affirming its role as a repeatable and valuable strategic tool not just for business ecosystem analysis, but also for other strategic decision-making contexts.

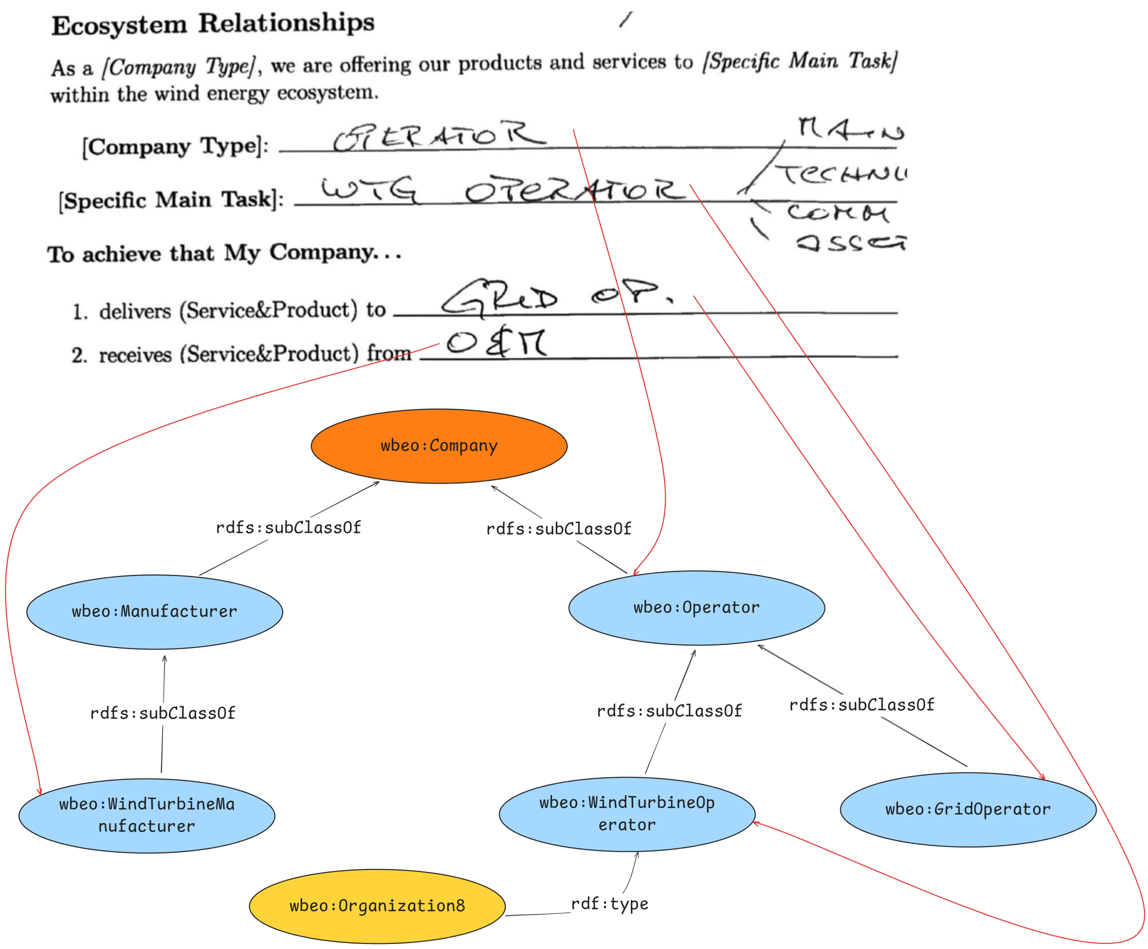

Figure 2: The BEAR framework transforms semi-structured survey data into a formal Application Ontology. This figure shows the part of the Value Creation process, where natural language responses from a business survey filled by Organization8 (top) are mapped through RU analysis into the explicit semantic relationships of an Application Ontology (bottom). Yellow represents the individual Organization8, blue represents the extended classes , and red represents the seed ontology entity. For the structure of the survey, see [11].

5. Discussion

Uncovering complex structural relationships and strategic blind spots within business ecosystems demands a semantic approach, not just syntactic analysis. While existing OKGE methodologies are inherently semantic, they do not explicitly organize and focus their engineering processes based on business goals, their jobs, and knowledge gaps of the stakeholders [7, 37, 38, 23, 34, 24, 21]. BEAR is engineered precisely to bridge this gap, inspired by well established OKGE methodologies, anchoring them within a value-first paradigm tailored for a single purpose: to answer what a decision-maker needs to know to get the job done [27].

To answer these KQs effectively, BEAR's robust handling of incomplete data—an everyday reality in business ecosystems yet unaddressed in literature [1, 12, 14, 15]—is a key capability for uncovering blind spots. For example, consider interactions occurring with a type of entity rather than a specific named individual. When data indicates an interaction with an unspecified entity (e.g, wbeo:Organization11 wbeo:deliversTo "some" engineering consultant company), BEAR models this target using reification; [rdf:type wbeo:EngineeringConsultantCompany]. This semantic modeling enables DL reasoners to deduce implicit connections and consequently reveals implicit connections and reveals blind spots—like consulting companies playing intermediary roles (Figure 4), that would otherwise remain hidden in traditional business ecosystem analysis.

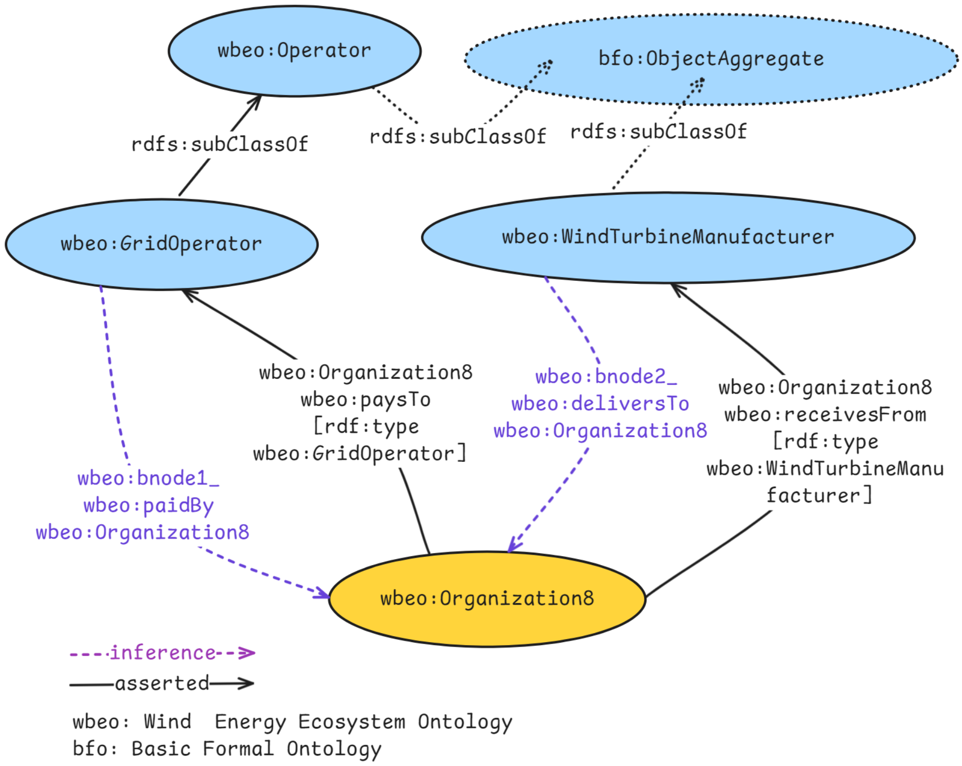

Figure 3: How BEAR Uses Reification to Connect Incomplete Data. The black asserted edges show the raw data: Organization8 (owl:NamedIndividual) pays to an unnamed Grid Operator (bnode1) and receives some product or service from some Wind Turbine Manufacturer (bnode2). Through inverse object property and reification, the DL reasoner infers the implicit relationships (purple dashed line): Some Grid Operator is getting paid by the Organization8, and some Wind Turbine Manufacturer is delivering some product or service to Organization8. This semantic inference is the engine that connects disparate data points into a coherent knowledge graph. Note: To demonstrate the reification more explicitly, we have shown one additional relation, which belongs to another knowledge question about the cash flows. The dotted bfo:ObjectAggregate entity shows a path for future work: aligning the model with an upper-level ontology [28].

5.1. The Modularity Paradox: Seeds, Silos, and Scale

Our seed ontology approach presents a fundamental tension in ontology engineering. By designing minimal ontologies for specific KQs, we achieve direct value delivery where traditional methodologies struggle. Yet, this focus risks creating "conceptual silos"—isolated ontologies that answer one question with pragmatic sufficiency, but do not communicate.

Our solution to this paradox lies in BEAR's relationship to modular ontology engineering philosophy. Our seed ontologies are inherently modular, not in their structure, but in their value, a different sense of MOMo [26]. Each seed ontology is a self-contained "value module", that delivers specific insight. The architectural challenge is, however, to prevent these modules from becoming disconnected. How do we prevent different knowledge questions from leading to disconnected seed ontologies? For example, in our pilot, we answered in total of 15 distinct KQs, and while we could essentially merge them into a single Wind Business Ecosystem Ontology (WBEO) [11], a systematic integration of these seed ontologies are needed to prevent siloing, especially as the domain of the KQs change significantly: We had to answer KQs about the different flows, at the same time, look at the importance of the operational data within it (See Survey [11]).

Our preliminary answer lies in treating seed ontologies as specialized modules in an evolving reference ontology. Evidently, the WBEO began from a single KQ; however, with iterative refinement—each new data point, each emergent class (like wbeo:Operator abstraction)—it moved closer to a comprehensive domain model (See Figure 2). This suggests a clear development pathway: seed ontologies for immediate value, which iteratively build and enrich a larger reference ontology for a long-term knowledge asset.

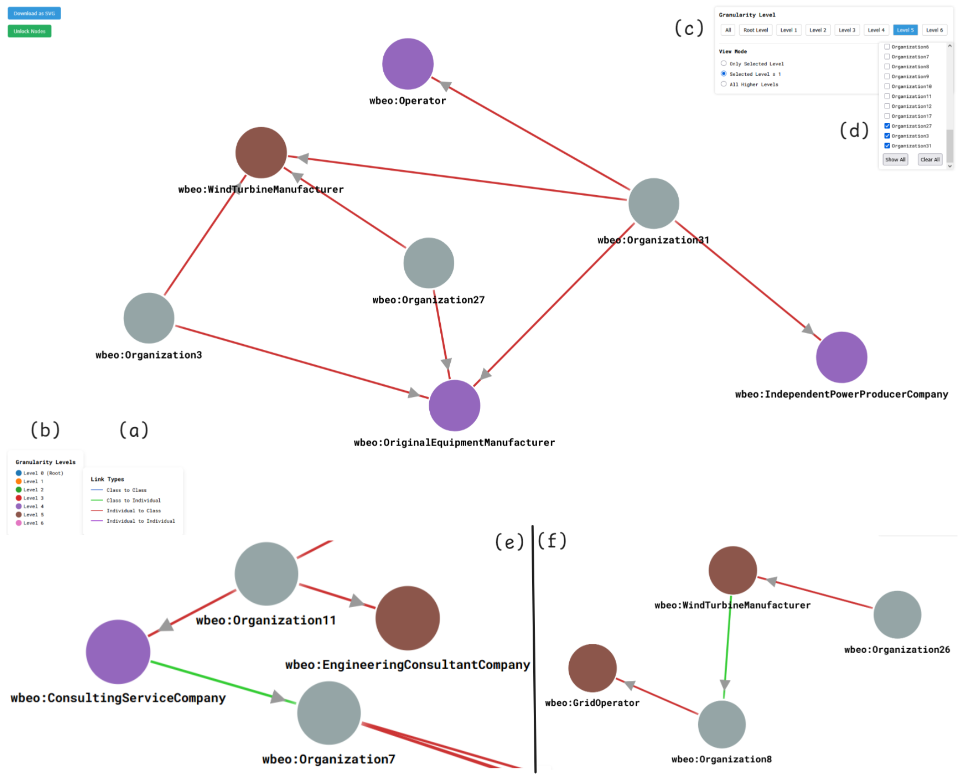

Figure 4: Interactive visualization uncovers hidden intermediary roles in the wind energy ecosystem. This tailored visualization, the final output of the BEAR framework, answers a stakeholder's knowledge question by mapping the complex ecosystem. Interactive features, such as (c) granularity adjustment and (d) named individual filtering, enable the discovery of critical blind spots while answering the given knowledge question. For instance, the graph exposes (e) hidden intermediary roles of consulting service companies and reveals that (f) Organization8 acts as a key bridge between wind turbine manufacturers and grid operators. Visual conventions aid interpretation for the decision makers: (a) red edges show individual-to-class relationships, while green edges show class-to-individual relationships. (b) Node colours represent the granularity level, with gray being the named individuals.

However, we argue that to achieve this, we must integrate an upper ontology (e.g., Basic Formal Ontology [28]) to map each seed ontology to a common framework.

5.2. Quality Metrics for Value-First Ontology Engineering

A value-first paradigm demands value-first metrics. While foundational for model soundness, traditional metrics like technical correctness and completeness measure the engineering artifact, not its business impact. They can tell us if a model is built correctly (e.g., by validation through competency questions [7, 23, 24]), but not if we built the right thing for the organizations.

We argue that, to bridge this gap, a project's success must be judged by its disposition to fill the stakeholders' knowledge gap. Based on the insights from our pilot project, we propose an initial set of value-first quality metrics (Table 4). This initial set permits further research and systematic validation.

5.3. Limitations and Future Work

Our framework, while promising, evidently has apparent limitations that define our future work. The current implementation relies on manual data collection and RU analysis, a critical scalability concern raised by our pilot stakeholders, which is our most explicit challenge. While our semi-structured survey is a step toward rapid, ontology-aligned data acquisition at events like industrial expos, true scalability requires automation. Therefore, we have initiated a new pilot focused on automated blind spot detection using graph pattern recognition. We are also exploring the use of LLMs to semi-automate RU extraction

Table 4
Value-First Quality Metrics for Ontology Engineering.

Metric	Guiding Question	Rationale: Why it Matters
Strategic Alignment	Do the developed ontologies, collected data, knowledge graph, etc., directly address the stakeholders' core business problem?	Ensures the work is aligned with strategic priorities, not just technical curiosity.
Actionability	Does the final answer enable a specific, informed decision or action?	Bridges the gap between data and execution; the output must be usable.
Pragmatic Sufficiency	Is the developed ontology sufficiently lean to deliver the maximum insight required by the knowledge question?	Prevents over-engineering and focuses resources on value, not exhaustive modeling.
Time-to-Value	How quickly does the engineering process move from business goal to actionable insight?	Measures the agility of the framework, which is critical for decision-makers on tight timelines.
Knowledge Gap Closure	To what degree was the stakeholder's initial knowledge gap resolved?	The ultimate success metric, rooted in the Jobs-To-Be-Done framework [3].

from unstructured sources, using a human-in-the-loop mechanism to accelerate the RU analysis and mapping process, directly aiming to improve the time-to-value quality metric (See Table 4) [39].

Another important aspect is the philosophical grounding of the BEAR framework. Our implicit commitment to scientific and ontological realism [40, 41], unlike cognitive [42] or linguistic [43], which aligns with the Basic Formal Ontology (BFO) [28], must be formalized to provide a more solid grounding for future development and integration.

Finally, while our single case study demonstrates feasibility and the successful handling of 15 distinct KQs suggests adaptability, generalizability must be validated. Even though it is due to its focus on one ecosystem (wind energy) and a subset of data (37 surveys), we have already begun a new pilot project in a different business context to test BEAR's applicability across diverse contexts and stakeholders, and we will share these results in future work.

As a conclusion, the complexity of modern business ecosystems demands semantic analysis—this much has been shown in our wind energy study. However, semantic capability alone is insufficient. Strategic decision makers do not need ontologies; they need answers. BEAR demonstrated this by inverting the traditional engineering process—starting with the answer they needed rather than the model required—we can deliver immediate strategic value while building toward a comprehensive knowledge infrastructure. The question is no longer whether ontologies can support strategic decision making, but whether we are willing to reorganize our engineering practices around the value they must deliver.

Acknowledgement

The financial support by the Austrian Federal Ministry of Digital and Economic Affairs, the National Foundation for Research, Technology, and Development, and the Christian Doppler Research Association is gratefully acknowledged.

Declaration on Generative AI

The authors used Grammarly for grammar and spelling support. The final manuscript was reviewed and edited by the authors, who are fully responsible for its content.

References

[1] R. Adner, Ecosystem as structure: An actionable construct for strategy, Journal of management 43 (2017) 39–58.

[2] J. F. Moore, Predators and prey: a new ecology of competition, Harvard business review 71 (1993) 75–86.

[3] C. M. Christensen, T. Hall, K. Dillon, D. S. Duncan, Know your customers' jobs to be done, Harvard business review 94 (2016) 54–62.

[4] R. Arista, X. Zheng, J. Lu, F. Mas, An ontology-based engineering system to support aircraft manufacturing system design, Journal of Manufacturing Systems (2023).

[5] A. Tüzün, S. Tripathi, N. Bachmann, A.-K. Thienemann, M. Brunner, H. Jodlbauer, Granular and relational swot analysis: An ontological approach, Procedia Computer Science 253 (2025) 2575–2585. URL: https://www.sciencedirect.com/science/article/pii/S1877050925003254. doi:https://doi.org/10.1016/j.procs.2025.01.317, 6th International Conference on Industry 4.0 and Smart Manufacturing.

[6] B. Smith, J. Williams, S.-K. Steffen, The ontology of the gene ontology, in: AMIA Annual Symposium Proceedings, volume 2003, 2003, p. 609.

[7] M. Fernández-López, A. Gómez-Pérez, N. Juristo Juzgado, Methontology: from ontological art towards ontological engineering (1997).

[8] C. Acciarini, F. Brunetta, P. Boccardelli, Cognitive biases and decision-making strategies in times of change: a systematic literature review, Management Decision 59 (2020) 638–652. URL: https://doi.org/10.1108/MD-07-2019-1006. doi:10.1108/MD-07-2019-1006.

[9] L. Pratt, C. Bisson, T. Warin, Bringing advanced technology to strategic decision-making: The decision intelligence/data science (di/ds) integration framework, Futures 152 (2023) 103217. URL: https://www.sciencedirect.com/science/article/pii/S0016328723001222. doi:https://doi.org/10.1016/j.futures.2023.103217.

[10] I. Farida, D. Setiawan, Business strategies and competitive advantage: The role of performance and innovation, Journal of Open Innovation: Technology, Market, and Complexity (2022). doi:10.3390/joitmc8030163.

[11] A. Tüzün, BEAR, https://github.com/T-Z-N/BEAR, 2025.

[12] Complex business ecosystem intelligence using ai-powered visual analytics, Decision Support Systems 178 (2024) 114133. doi:https://doi.org/10.1016/j.dss.2023.114133.

[13] Unlocking hidden market segments: A data-driven approach exemplified by the electric vehicle market, Expert Systems with Applications 254 (2024) 124331. doi:https://doi.org/10.1016/j.eswa.2024.124331.

[14] R. C. Basole, T. Clear, M. Hu, H. Mehrotra, J. Stasko, Understanding interfirm relationships in business ecosystems with interactive visualization, IEEE Transactions on Visualization and Computer Graphics 19 (2013) 2526–2535. doi:10.1109/TVCG.2013.209.

[15] R. Basole, M. Russell, J. Huhtamäki, N. Rubens, K. Still, H. Park, Understanding business ecosystem dynamics: A data-driven approach, ACM Transactions on Management Information Systems 6 (2015). doi:10.1145/2724730.

[16] C. H. Tian, B. K. Ray, J. Lee, R. Cao, W. Ding, Beam: A framework for business ecosystem analysis and modeling, IBM Systems Journal 47 (2008) 101–114. doi:10.1147/sj.471.0101.

[17] S. Rinkinen, V. Harmaakorpi, The business ecosystem concept in innovation policy context: building a conceptual framework, Innovation: The European Journal of Social Science Research 31 (2018) 333–349. doi:10.1080/13511610.2017.1300089.

[18] J. Gordijn, E-business value modelling using the e3-value ontology, in: Value creation from e-business models, Elsevier, 2004, pp. 98–127.

[19] N. Bartels, R. Wieringa, M. Koch, K. Villela, D. Suzumura, J. Gordijn, Using tangible modeling to create an e3 value conceptual model for digital ecosystems, in: IFIP Working Conference on The Practice of Enterprise Modeling, Springer, 2022, pp. 134–148.

[20] S. Staab, R. Studer, H.-P. Schnurr, Y. Sure, Knowledge processes and ontologies, IEEE Intelligent

systems 16 (2001) 26–34.
[21] A. De Nicola, M. Missikoff, A lightweight methodology for rapid ontology engineering, Communications of the ACM 59 (2016) 79–86.
[22] A. De Nicola, M. Missikoff, R. Navigli, A proposal for a unified process for ontology building: Upon, in: International conference on database and expert systems applications, Springer, 2005, pp. 655–664.
[23] M. Poveda-Villalón, A. Fernández-Izquierdo, M. Fernández-López, R. García-Castro, Lot: An industrial oriented ontology engineering framework, Engineering Applications of Artificial Intelligence 111 (2022) 104755.
[24] M. Hristozova, L. Sterling, An extreme method for developing lightweight ontologies, in: Workshop on Ontologies in Agent Systems, 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems, 2002.
[25] M. Boicu, G. Tecuci, D. Marcu, Ontologies and the Knowledge Acquisition Bottleneck, in: Proceedings of the IJCAI-01 Workshop on Ontologies and Information Sharing, Seattle, WA, USA, 2001, pp. 12–18.
[26] C. Shimizu, K. Hammar, P. Hitzler, Modular ontology modeling, Semantic Web 14 (2023) 459–489.
[27] C. Christensen, T. Hall, K. Dillon, D. S. Duncan, Competing against luck, The story of innovation and customer choice. First edition. New York, NY: HarperBusiness an imprint of HarperCollins Publishers (2016).
[28] R. Arp, B. Smith, A. D. Spear, Building ontologies with basic formal ontology, Mit Press, 2015.
[29] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, S. Rudolph, et al., Owl 2 web ontology language primer, W3C recommendation 27 (2009) 123.
[30] M. Giunti, G. Sergioli, G. Vivanet, S. Pinna, Representing n-ary relations in the semantic web, Logic Journal of the IGPL 29 (2021) 697–717.
[31] M. Bostock, D3, https://github.com/d3/d3, 2025. [Online; accessed 2025-07-10].
[32] S. Lohmann, V. Link, E. Marbach, S. Negru, Webvowl: Web-based visualization of ontologies, in: International Conference on Knowledge Engineering and Knowledge Management, Springer, 2014, pp. 154–158.
[33] Graphdb, https://graphdb.ontotext.com/documentation, 2025. [Online; accessed 2025-08-02].
[34] M. Uschold, M. Gruninger, Ontologies: principles, methods and applications, The Knowledge Engineering Review 11 (1996) 93–136. doi:10.1017/S0269888900007797.
[35] WindEnergy Hamburg, Global on & offshore event, 2024. URL: https://www.windenergyhamburg.de/, [Online; accessed 2025-05-10].
[36] D. Küçük, D. Küçük, Ontowind: An improved and extended wind energy ontology, arXiv preprint arXiv:1803.02808 (2018).
[37] D. Vrandečić, S. Pinto, C. Tempich, Y. Sure, The diligent knowledge processes, Journal of Knowledge Management 9 (2005) 85–96.
[38] A. Gómez-Pérez, M. C. Suárez-Figueroa, Neon methodology for building ontology networks: a scenario-based methodology (2009).
[39] C. Shimizu, P. Hitzler, Accelerating knowledge graph and ontology engineering with large language models, Journal of Web Semantics (2025) 100862.
[40] A. Chakravartty, Scientific Realism, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, summer 2017 ed., Metaphysics Research Lab, Stanford University, 2017. URL: https://plato.stanford.edu/archives/sum2017/entries/scientific-realism/.
[41] B. Smith, W. Ceusters, Ontological realism: A methodology for coordinated evolution of scientific ontologies, Applied ontology 5 (2010) 139–188.
[42] G. Guizzardi, A. Botti Benevides, C. M. Fonseca, D. Porello, J. P. A. Almeida, T. Prince Sales, Ufo: Unified foundational ontology, Applied ontology 17 (2022) 167–210.
[43] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, L. Schneider, Sweetening ontologies with dolce, Springer, 2002, pp. 166–181.