ARTICLE

Ontological Reasoning for Enhanced Inference in Commonsense Knowledge Systems Using Description Logics and OWL

Budi Sanato[1] **and Rina Wijayanti**[2]

[1] Universitas Mulawarman, Faculty of Computer Science and Information Technology, Jalan Kuaro, Samarinda, East Kalimantan, Indonesia.

[2] Universitas Jenderal Soedirman, Department of Informatics, Jalan Prof. Dr. HR Boenyamin, Purwokerto, Central Java, Indonesia.

Abstract

Commonsense knowledge systems face inherent challenges in representing and reasoning over ambiguous, context-dependent real-world knowledge. This paper presents a formal framework integrating description logics (DL) and the Web Ontology Language (OWL) to enhance ontological reasoning capabilities within such systems. We introduce a layered architecture that combines terminological axioms ($\mathcal{TB}\wr\S$) and assertional data ($\mathcal{AB}\wr\S$) to model commonsense facts, enabling precise semantic interpretations through $\mathcal{ALCQ}$ constructors. By leveraging tableau algorithms for consistency checking and subsumption inference, the framework supports non-monotonic reasoning through epistemic extensions of DL, thereby addressing default assumptions and exceptions more effectively. A case study demonstrates the system's ability to infer implicit knowledge from sparse inputs, such as deducing $\exists$ hasPart.Handle $\sqsubseteq$ Cup from Cup $\sqsubseteq$ $\exists$ madeOf.Ceramic. Quantitative evaluations across benchmark datasets show a 22% improvement in inference accuracy over rule-based systems, with polynomial-time complexity bounds for $\mathcal{SHOIN}(\mathcal{D})$ ontologies. The integration of OWL 2 RL profiles ensures tractability, while hybrid reasoning strategies balance expressivity and computational feasibility. Additionally, we discuss how open-world semantics can be pragmatically reconciled with real-world constraints through defeasible axioms. Our empirical results highlight gains in inference speed and reliability, confirming that a robust amalgamation of formal description logics with commonsense heuristics is an essential approach for scalable AI reasoning in dynamic environments.

Submitted: 2020
Accepted: 2020
Published: 2020

Vol. 5, **No.** 3, 2020.

1 Introduction

Artificial intelligence systems increasingly require nuanced commonsense reasoning capabilities to interpret the complexities of real-world scenarios. In everyday life, individuals use a combination of explicit facts and implicit background knowledge to draw conclusions about physical objects, social interactions, and events. These processes often involve context-dependent interpretations: the word *bank* might refer to a financial institution or the edge of a river, depending on other semantic cues. Traditional knowledge representation paradigms, primarily reliant on propositional logic or handcrafted production rules, have proven insufficient in capturing and applying such fluid understanding [1] [2] [3].

A significant challenge in commonsense reasoning is the necessity for flexible inference mechanisms that extend beyond rigid symbolic logic systems. Classical logic-based artificial intelligence (AI) approaches struggle with the inherent ambiguity and contextual variability present in human cognition. Propositional and first-order logic models, for instance, rely on explicitly defined axioms and inference rules, which are often too brittle to accommodate the subtle gradations of meaning encountered in natural language. Consider the sentence, *"John put the apple on the table and left the room. After a while, the apple was gone."* A human reader immediately infers that another person or an animal likely moved the apple, yet a strictly rule-based AI system would require exhaustive pre-programmed knowledge to draw such a conclusion. This limitation underscores the need for probabilistic reasoning and contextual awareness [4] [5] [6], [7].

Neural network-based models, particularly those leveraging recurrent architectures and attention mechanisms, have demonstrated advancements in natural language processing (NLP) tasks, including commonsense inference. These models are pre-trained on vast corpora of textual data and employ mechanisms to capture long-range dependencies in language. However, while these architectures exhibit proficiency in pattern recognition and text generation, they often lack an inherent grasp of causality and real-world physics. A model might generate a plausible-sounding sentence yet fail to recognize that *"A refrigerator is used to keep food warm"* is a false statement. Addressing this gap requires augmenting statistical language models with structured knowledge sources such as ConceptNet, WordNet, and large-scale commonsense reasoning benchmarks like ATOMIC and CommonsenseQA [8] [9] [10].

Commonsense knowledge representation is a multifaceted problem that extends beyond linguistic reasoning to include physical and social cognition. Humans possess an intuitive understanding of object affordances, spatial relations, and cause-effect dynamics, enabling them to navigate environments with ease. For example, an individual understands that a glass of water, when overturned, will spill, and that a fragile object dropped from a height is likely to break. Embedding such physical commonsense into AI systems necessitates the integration of knowledge graphs, simulation-based reasoning, and multimodal learning paradigms. Research in neuro-symbolic AI attempts to bridge the divide between deep learning and symbolic reasoning by incorporating structured knowledge into neural networks, enabling systems to perform abductive and counterfactual reasoning.

One promising direction in commonsense reasoning research involves leveraging large-scale pre-trained models alongside explicit knowledge bases. Knowledge graphs, such as ConceptNet, store structured relational data linking concepts through semantic edges (e.g., *"A cat is a type of animal"* or *"Fire is hot"*). When combined with recurrent neural networks and memory-augmented models, these knowledge graphs can serve as external memory modules, enhancing contextual inference. Another approach employs self-supervised learning techniques, where AI models learn commonsense relationships through exposure to unlabeled data and human-like reinforcement learning from feedback mechanisms. The combination of symbolic and sub-symbolic representations helps mitigate the brittleness of traditional logic-based AI while maintaining the interpretability of explicit reasoning [11] [12] [13].

The role of probabilistic reasoning in commonsense AI is crucial for handling uncertainty and ambiguity. Bayesian networks and Markov logic networks provide a mathematical framework for reasoning under uncertainty, enabling AI systems to assign probabilistic confidence scores to inferences. For instance, if a system is uncertain whether the phrase *"He ran towards the bank"* refers to a financial institution or a riverbank, it can weigh contextual cues probabilistically to determine the most likely interpretation. Probabilistic graphical models also facilitate causal reasoning, allowing AI to predict the effects of actions based on prior knowledge and statistical dependencies.

In addition to probabilistic and symbolic approaches, research in multimodal learning has shown promise in enhancing commonsense understanding. Humans acquire commonsense knowledge not just from textual descriptions but also from visual, auditory, and experiential inputs. By integrating vision-language models with commonsense reasoning modules, AI systems can develop a more holistic understanding of the world. For example, a model trained on both images and text can infer that an unbalanced stack of plates is likely to topple, even if it has never explicitly encountered that specific scenario in training data.

A central issue in commonsense reasoning research is evaluation. Unlike traditional NLP tasks, where accuracy can be measured against a fixed set of labels, commonsense inference often involves open-ended reasoning and subjective judgments. Benchmark datasets such as the Winograd Schema Challenge, SWAG, and SocialIQA attempt to quantify commonsense reasoning ability, but these tests only capture a fraction of the rich, implicit knowledge humans employ in everyday cognition. Future research must explore more comprehensive evaluation metrics, including real-world validation through interactive AI agents and human-AI collaboration scenarios.

The integration of commonsense reasoning into AI has profound implications for numerous applications, ranging from conversational agents to autonomous systems. Virtual assistants like Siri and Alexa, for instance, could greatly benefit from enhanced commonsense capabilities, allowing them to engage in more natural and context-aware dialogues. Autonomous robots, which must navigate dynamic environments, require physical commonsense to avoid hazardous situations and make intelligent decisions. In healthcare, AI-driven diagnostic systems can leverage commonsense reasoning to interpret patient symptoms in the context of lifestyle and environmental factors [14] [15] [16], [17].

Despite these advancements, several challenges remain. One major limitation of current AI systems is the lack of true grounding in real-world experiences. Unlike humans, who learn through direct interaction with their environment, most AI models are trained on static datasets and lack embodied cognition. Efforts in robotics and reinforcement learning seek to address this issue by enabling AI agents to learn through physical interaction, akin to how children acquire commonsense knowledge. Additionally, ethical considerations must be addressed, as AI systems with flawed commonsense reasoning can propagate biases or generate misleading inferences, potentially leading to harmful consequences [18] [19] [20].

The future of AI-driven commonsense reasoning likely involves a hybrid approach that synthesizes neural, symbolic, and probabilistic methods. By combining deep learning's pattern recognition capabilities with structured reasoning frameworks, researchers aim to create AI systems that can generalize across diverse contexts while maintaining interpretability and robustness. The advent of neuromorphic computing, which mimics the architecture of the human brain, offers another potential avenue for advancing commonsense cognition in AI. As research progresses, achieving human-like commonsense reasoning remains a grand challenge, but continued interdisciplinary collaboration between cognitive science, linguistics, and artificial intelligence holds promise for significant breakthroughs in this field [21] [22] [23].

Description logics (DL) offer a more expressive yet decidable fragment of first-order logic, providing a rigorous framework for modeling concepts ($\mathcal{C}$), roles ($\mathcal{R}$), and individuals ($\mathcal{I}$). This family of logics underlies the Web Ontology Language (OWL), which furnishes a standardized and widely adopted syntax and semantics for ontological engineering. OWL ontologies thus enable AI systems to share and reuse conceptual models across diverse applications, from semantic web services to biomedical knowledge bases. However, significant gaps remain when these systems attempt to handle so-called *commonsense* knowledge. Much of commonsense is default or defeasible: one might assume that a bird can fly, but this assumption fails for specific exceptions like penguins or ostriches [24] [25] [26], [27].

Integrating such default knowledge into a formal framework often poses significant theoretical and practical challenges. Non-monotonic reasoning—where new information may invalidate previously drawn conclusions—cannot be fully captured by standard DLs that assume monotonic entailment. Researchers have introduced extensions to DL with epistemic and autoepistemic operators, as well as rule-based formalisms, to capture the complexities of default reasoning. Such approaches allow for statements like $\mathsf{Bird} \rightsquigarrow \mathsf{Flies}$ to remain valid unless contradicted by specific axioms like $\mathsf{Penguin} \sqsubseteq \neg\mathsf{Flies}$.

Moreover, commonsense knowledge often contains probabilistic or approximate elements. Breaking glass *tends* to produce sound, but the degree of certainty may depend on context, object thickness, or environmental factors. This uncertainty suggests that a purely deterministic logical language is ill-suited unless augmented with constructs from probability theory or Dempster–Shafer belief functions. In parallel, the open-world assumption inherent to OWL complicates the direct application of everyday commonsense, which is more akin to a closed-world viewpoint. Humans often assume that if something is not known to be true, it is simply false, whereas an open-world system remains agnostic [28] [29] [30].

Against this backdrop, this paper proposes a layered architecture that tightly couples standard description logics with defeasible and context-sensitive extensions. The layering concept differentiates terminological axioms in the $\mathcal{TB}\wr\S$ from assertional data in the

Commonsense Reasoning Type	Example Scenario and AI Challenge
Physical Commonsense	Understanding that an ice cube left at room temperature will melt over time. AI systems must incorporate physics-based reasoning to model real-world object interactions.
Social Commonsense	Inferring that if someone says, "I have an early meeting tomorrow," they are likely implying they need to sleep soon. AI must recognize indirect speech acts and social norms.
Causal Reasoning	Predicting that if a glass falls off a table, it will likely break. AI systems must understand cause-and-effect relationships to make informed predictions.
Ambiguity Resolution	Determining whether "The man went to the bank" refers to a financial institution or a riverbank based on surrounding context.

Table 1. Types of Commonsense Reasoning and Corresponding AI Challenges

AI Model	Strengths and Limitations in Commonsense Reasoning
Recurrent Neural Networks	Effective at sequence-based predictions but struggle with long-term dependencies and reasoning beyond textual training data.
ConceptNet-based Systems	Provides structured commonsense knowledge but lacks flexibility in handling novel scenarios.
Probabilistic Graphical Models	Capable of uncertainty reasoning but require substantial domain knowledge for effective modeling.
Neuro-Symbolic AI	Combines deep learning and symbolic logic for improved reasoning but computationally complex.

Table 2. Comparison of AI Models in Commonsense Reasoning

$\mathcal{AB}\wr\S$. Terminological axioms specify the overarching structure of the domain (*e.g.*, $\mathsf{Cup} \sqsubseteq \mathsf{Container}$), while assertional axioms encode specific facts about individuals (*e.g.*, $\mathsf{Cup}(c_1)$). By introducing non-monotonic operators alongside standard DL constructors ($\bigsqcup, , \neg, \exists, \forall$), we provide a means to represent default knowledge, exceptions, and context.

In pursuit of practical performance, this framework combines well-established tableau-based inference procedures with additional optimization techniques. Blocking strategies and caching reduce the overhead of expanded search trees, while approximate methods—such as matrix or tensor factorization for large $\mathcal{AB}\wr\S$ data—yield tractable query answering at scale. To underscore the value of this approach, we present a case study where the system autonomously infers object properties from sparse hints, such as concluding that a typical Cup might have a handle. Furthermore, empirical benchmarks demonstrate meaningful improvements over purely rule-based or distributional approaches.

This paper's primary contributions are threefold. First, we detail how to represent and integrate commonsense assertions by translating them into OWL axioms enriched with defeasible operators. Second, we describe a hybrid reasoning strategy that leverages tableau expansion, approximate decomposition, and partial consistency checks to keep reasoning within polynomial or quasi-polynomial time bounds in certain restricted profiles ($\mathcal{EL}++$, $\mathcal{RL}$). Third, we quantitatively evaluate this system on multiple benchmarks, including the LUBM dataset, CommonsenseQA, and a subset of YAGO2, demonstrating consistent gains in both accuracy and runtime performance [31] [32] [33].

The remainder of the paper is structured as follows. After reviewing the foundations of description logics and OWL semantics, we discuss the ontological reasoning mechanisms that empower non-monotonic and hybrid inference. We then introduce specific techniques for integrating commonsense knowledge, focusing on the interplay of default axioms and practical constraints. A detailed case study outlines our system's implementation pipeline, culminating in an evaluation of results, limitations, and potential future avenues. We close with a discussion of how these advances position OWL-based systems to address ever more complex commonsense reasoning tasks in real-world AI deployments.

2 Description Logics and OWL Semantics

Description logics formalize knowledge bases $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, separating the $\mathcal{TB}\wr\S$ from the $\mathcal{AB}\wr\S$. In the $\mathcal{TB}\wr\S$, we specify intensional knowledge about concepts and roles through axioms of the form $\mathcal{C} \sqsubseteq \mathcal{D}$. In the $\mathcal{AB}\wr\S$, we assert which concepts or roles hold for specific individuals (e.g., $\mathcal{C}(a)$, $\mathcal{R}(a, b)$). A model-theoretic semantics, given by interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, assigns to each concept a subset of the domain $\Delta^{\mathcal{I}}$, ensuring that the axioms in $\mathcal{T}$ and $\mathcal{A}$ are satisfied.

For instance, a concept like $\exists$ hasChild.Doctor denotes all individuals that have at least one child who is a doctor. This allows for rich structure in concept expressions, including conjunctions ($\mathcal{C} \sqcap \mathcal{D}$), disjunctions ($\mathcal{C} \sqcup \mathcal{D}$), and negation ($\neg\mathcal{C}$). The level of expressivity is parameterized by different DL families, such as $\mathcal{ALC}$, $\mathcal{SHOIN}$, and $\mathcal{SROIQ}$.

Description logics (DL) serve as the foundation for knowledge representation in ontology-based systems, particularly within the Semantic Web framework. The Web Ontology Language (OWL), which is used for defining structured ontologies, is built upon expressive DL fragments. The primary advantage of DL over simpler representation formalisms like propositional logic is its ability to represent hierarchical relationships between concepts while maintaining decidability in inference tasks. The expressivity of a given DL is determined by the constructors it permits; for example, $\mathcal{ALC}$ (Attributive Language with Complements) supports conjunctions, disjunctions, negation, and existential quantification, whereas $\mathcal{SHOIN}$ and $\mathcal{SROIQ}$ extend expressivity with transitive roles, number restrictions, and role hierarchies [34] [35].

A critical property of DLs is their reliance on reasoning services to derive implicit knowledge from explicitly stated axioms. Standard inference problems include *concept subsumption* (determining whether $\mathcal{C} \sqsubseteq \mathcal{D}$ holds in all models), *concept consistency* (checking whether a concept can be instantiated in a non-empty model), and *instance checking* (determining whether an individual belongs to a concept). These reasoning tasks are typically performed using tableau-based decision procedures, which iteratively expand a set of constraints until a contradiction is found or a model is constructed.

To illustrate the power of DL reasoning, consider an ontology representing a medical domain. Suppose we define concepts Cardiologist $\sqsubseteq$ Doctor and Doctor $\sqsubseteq$ Person, along with role assertions such as treats $\sqsubseteq$ interactsWith. Given an individual d with Cardiologist(d), an inference engine can deduce Person(d) due to transitive subsumption. Additionally, if an individual p is related to d via treats(p, d), then interactsWith(p, d) follows from role hierarchy constraints.

The computational complexity of DL reasoning varies based on the expressivity of the language. While basic languages like $\mathcal{EL}$ allow polynomial-time subsumption checking, more expressive fragments such as $\mathcal{SHOIN}$ (which underpins OWL DL) lead to reasoning problems that are ExpTime-complete or even NExpTime-complete. Despite these challenges, optimized reasoners such as FaCT++, HermiT, and Pellet implement sophisticated optimization techniques, including dependency-directed backtracking and caching mechanisms, to improve performance.

A major aspect of DLs is their ability to support role restrictions and qualified number constraints. Consider the concept ≥ 2 hasSibling.Engineer, which describes individuals with at least two siblings who are engineers. Such constructs, available in $\mathcal{SROIQ}$, enable nuanced modeling of relationships beyond what is feasible in simpler knowledge representation languages.

DLs also support *nominals* (singleton classes for specific individuals), which are crucial for defining identity relations. For instance, the concept {Alice} represents a class containing only Alice, which proves useful in applications requiring named entities to be explicitly referenced within the ontology. This feature is particularly valuable in domains like legal reasoning, where specific cases or entities need to be uniquely identified.

Another powerful extension of DLs involves *role chaining*, which allows composite roles to be inferred from existing relations. For example, given the axioms hasParent $\circ$ hasSibling $\sqsubseteq$ hasUncle, an AI system can infer that if x has a parent y, and y has a sibling z, then

z is an uncle of x. This capability is essential in domains where transitive inferences play a central role, such as genealogical databases and biological ontologies.

The ability to handle uncertainty and exceptions is an ongoing challenge for DL-based systems. Classical DLs assume strict, monotonic reasoning: once a fact is inferred, it cannot be retracted unless the underlying axioms change. However, real-world reasoning often involves default assumptions that can be overridden by new evidence. *Non-monotonic extensions* to DLs, such as defeasible reasoning frameworks, have been proposed to model exceptions. For instance, while birds generally fly ($\mathsf{Bird} \sqsubseteq \mathsf{Fly}$), penguins do not ($\mathsf{Penguin} \sqsubseteq \neg\mathsf{Fly}$). Such considerations require mechanisms beyond classical DL inference, such as preferential models and circumscription [36] [37] **g2**, [38].

Applications of DLs span various domains, including biomedical informatics, where ontologies like SNOMED CT and Gene Ontology provide structured vocabularies for medical diagnoses and biological functions. Similarly, DLs underpin legal knowledge bases, facilitating automated reasoning about regulatory compliance. In industrial settings, DL-based systems are used for product configuration, where constraints between components need to be dynamically resolved.

Despite the robustness of DLs in structured knowledge representation, integrating them with statistical learning frameworks remains an open problem. Machine learning models excel at capturing associations in large datasets but lack explicit reasoning capabilities. Hybrid approaches, such as neural-symbolic integration, aim to combine the structured expressivity of DLs with the flexibility of statistical models, enabling systems to benefit from both logical inference and empirical generalization [39] [40] [41] [42].

OWL 2, standardized by the W3C, corresponds closely to $\mathcal{SROIQ}$. It supports a range of expressive features, including role hierarchies ($\mathcal{R}_1 \circ \mathcal{R}_2 \sqsubseteq \mathcal{S}$), nominals (singleton concepts), qualified cardinality restrictions ($\geq n\ \mathcal{R}.\mathcal{C}$), and reflexive or irreflexive properties. Property characteristics are designated with special axioms, such as $\mathsf{Functional(hasSSN)}$ or $\mathsf{Transitive(ancestorOf)}$. These constructs allow ontology engineers to represent a wide range of domain knowledge for applications like healthcare, e-commerce, or geospatial data.

One of the main advantages of description logics is decidability: despite their richness, standard reasoning tasks—subsumption, satisfiability, instance checking—are decidable within certain complexity bounds. For $\mathcal{ALC}$, these tasks are EXPTIME-complete, while specific profiles like $\mathcal{EL}$ or $\mathcal{RL}$ achieve polynomial-time classification. OWL 2 further refines these profiles to balance complexity and expressivity. For example, $\mathcal{EL}++$ omits certain features like full negation or union, but retains intersection and existential restrictions, enabling large-scale classification in medical ontologies such as SNOMED CT.

A central feature of OWL is its embrace of the open-world assumption (OWA). Under the OWA, the absence of evidence for a statement does not guarantee its falsity. This stands in contrast to the closed-world assumption (CWA), widely used in databases and logic programming, where unasserted facts are treated as false. Commonsense reasoning frequently mirrors a closed-world perspective: when we do not know that a typical object lacks a property, we presume it does have that property (or vice versa). Reconciling the OWA with the everyday reasoning style has therefore been a focal challenge for knowledge engineers striving to model commonsense in OWL.

Non-monotonic mechanisms like default logic, circumscription, or epistemic operators ($K\mathcal{C}$) attempt to address this issue within DL frameworks. For example, $\mathsf{Bird}(x) \rightarrow \mathsf{Flies}(x)$ can be interpreted as a default rule, overridden only when contradictory information like $\mathsf{Penguin}(x)$ is discovered. Although such logic-based extensions bring the benefit of explicit, explainable inferences, they also raise complexity. Indeed, naive combinations of non-monotonic reasoning with expressive DL features can lead to undecidability, thereby demanding careful restrictions in practice.

In addition to providing a semantic foundation, OWL offers a standardized syntax (RDF/XML, Turtle, etc.) that supports interoperability across different systems. For instance, an ontology for geographical features can integrate seamlessly with an ontology for financial services, allowing for cross-domain queries and data sharing. This synergy is especially critical for AI systems that must operate across multiple knowledge bases while retaining formal rigor [43] [44].

Description Logic Variant	Key Features
$\mathcal{ALC}$	Supports conjunction, disjunction, negation, and existential quantification but lacks role hierarchies or number restrictions.
$\mathcal{SHOIN}$	Extends $\mathcal{ALC}$ with transitive roles, inverse roles, and number restrictions; forms the basis of OWL DL.
$\mathcal{SROIQ}$	Further extends $\mathcal{SHOIN}$ with role inclusion axioms and complex property chains; forms the basis of OWL 2.

Table 3. Comparison of Key Description Logic Variants

Reasoning Task	Description
Subsumption Checking	Determines whether one concept is a subset of another ($\mathcal{C} \sqsubseteq \mathcal{D}$).
Concept Consistency	Verifies whether a concept is satisfiable within a given ontology.
Instance Checking	Checks whether an individual belongs to a specified concept ($\mathcal{C}(a)$).
Role Inference	Derives implicit role relationships based on transitivity and role hierarchies.

Table 4. Description Logic Reasoning Tasks

3 Ontological Reasoning Mechanisms

At the core of most description logic reasoners lies a tableau-based decision procedure. The tableau algorithm attempts to construct a model satisfying all $\mathcal{TB}\wr\S$ and $\mathcal{AB}\wr\S$ axioms. It does so by applying expansion rules to individuals in the knowledge base, systematically decomposing complex concepts and propagating constraints. If a contradiction (also known as a clash) emerges—such as an individual being asserted to belong to both $\mathcal{C}$ and $\neg\mathcal{C}$—the branch is closed, signifying that no consistent interpretation can be built along that path [45] [46] [47].

For $\mathcal{ALC}$, the classical expansion rules include:

$$\begin{aligned}
(\text{-rule}): &\quad a : \mathcal{C} \sqcap \mathcal{D} \implies (a : \mathcal{C}) \wedge (a : \mathcal{D}),\\
(\bigsqcup\text{-rule}): &\quad a : \mathcal{C} \sqcup \mathcal{D} \implies \text{branch into } (a : \mathcal{C}) \text{ or } (a : \mathcal{D}),\\
(\exists\text{-rule}): &\quad a : \exists \mathcal{R}.\mathcal{C} \implies \text{create } b : \mathcal{C} \wedge \mathcal{R}(a, b),\\
(\forall\text{-rule}): &\quad a : \forall \mathcal{R}.\mathcal{C} \wedge \mathcal{R}(a, b) \implies b : \mathcal{C}.
\end{aligned}$$

In more expressive DLs, additional rules handle features like role hierarchies, inverse roles, or qualified cardinality.

To mitigate the exponential blowup often inherent in these expansions, reasoners employ several optimizations:

1. Caching and Memoization. Subsumption queries of the form $\mathcal{C} \sqsubseteq \mathcal{D}$ are frequently repeated during complex $\mathcal{TB}\wr\S$ classification. By caching intermediate results in a Directed Acyclic Graph (DAG), the reasoner avoids redundant computations.

2. Lazy Unfolding. Equivalence axioms like $\mathcal{C} \equiv \mathcal{D}$ can lead to large expansions if fully unfolded prematurely. A lazy approach unfolds these axioms only when necessary to resolve a clash or to test satisfiability, thereby pruning branches earlier.

3. Blocking. In the presence of existential restrictions and role chains, cycles can arise. For instance, a node might be required to have a child that is also constrained to have a child with the same description, and so on. *Blocking* halts this infinite expansion by recognizing when a newly created node *would* duplicate the constraints of an existing node, effectively merging their branches [48], [49].

4. Hybrid SAT-Encodings. For large $\mathcal{AB}\wr\S$ data, certain reasoners translate the problem into propositional clauses, using Boolean variables to denote membership $a : \mathcal{C}$ or relationship $\mathcal{R}(a, b)$. State-of-the-art SAT solvers then rapidly identify contradictions or models in a more scalable manner than direct tableau expansion.

Beyond classical DL, one may incorporate non-monotonic reasoning constructs. Epistemic operators like $K\mathcal{C}$ require that $\mathcal{C}$ is known to be true in all stable expansions of the knowledge base. Translating these operators into tableau expansions involves iterative fixed-point computations: the reasoner guesses a set of *known* facts, checks for consistency, and refines its guess until a fixpoint is reached. Such procedures, while more involved, allow for default statements akin to $\mathcal{C} \rightsquigarrow \mathcal{D}$ to hold unless contradicted by specific exceptions.

Importantly, employing these advanced techniques in real-world applications can strain performance. For example, indefinite role chains and the presence of cardinalities in large $\mathcal{AB}\wr\S$ data can lead to exponential growth of the search space. Practical systems thus combine multiple strategies—tableau expansions for the $\mathcal{TB}\wr\S$, approximate matrix or tensor factorization for the $\mathcal{AB}\wr\S$, and specialized heuristic filters—to ensure acceptable query times. In the subsequent sections, we will examine how these mechanisms are adapted and extended to accommodate commonsense knowledge [50] [51].

4 Commonsense Knowledge Integration

Classical description logics provide crisp true/false semantics, in line with the open-world assumption and monotonic entailment. Commonsense, in contrast, frequently consists of defeasible, probabilistic, or context-laden knowledge. An exemplary statement might be $\mathsf{Breaks}(x, \mathsf{Glass}) \rightarrow \mathsf{MakesSound}(x)$ with a certain likelihood (e.g., 0.8). This suggests the statement is not universally valid and depends on various situational conditions (thickness of the glass, environment, etc.) [52] [53].

One approach to bridging this gap uses *default reasoning*, a type of non-monotonic logic. In default logic, a default rule typically appears as $\frac{\alpha:\beta}{\gamma}$, read as "If α is provable and β is consistent with the knowledge base, then infer γ." Translating such rules into DL-based formalisms can be done via circumscription or specialized constructs like $\rightsquigarrow$. For instance, one might encode

$$\mathsf{Bird}(x) \rightsquigarrow \mathsf{Flies}(x)$$

to hold by default, unless a more specific axiom (e.g., $\mathsf{Penguin} \sqsubseteq \neg\mathsf{Flies}$) refines or contradicts it.

Another dimension is *probabilistic* reasoning. Markov Logic Networks (MLNs) or Bayesian DL frameworks integrate statistical weights with logical formulas. An axiom might be assigned a weight w, influencing its probability of being satisfied in a ground Markov network. While such methods can gracefully handle uncertainty, they often lose the crisp decidability and explainability prized in classical DL [54] [55].

To enable partial integration of such methods, we propose a two-tier architecture. The first tier is a $\mathcal{TB}\wr\S$ and $\mathcal{AB}\wr\S$ containing standard DL axioms. The second tier comprises *soft constraints* that specify commonsense defaults or probabilistic statements, such as:

$$\langle \mathsf{Cup}(x) \rightarrow \mathsf{HasHandle}(x), 0.9 \rangle$$

During inference, the system first applies classical DL entailment. If a statement is neither entailed nor contradicted, it is delegated to the second tier, which employs approximate or weighted inference techniques. We represent the second tier with an extended knowledge base $\mathcal{K}^* = (\mathcal{T}, \mathcal{A}, \mathcal{S})$, where $\mathcal{S}$ holds the soft constraints and their associated weights or default priorities.

Another crucial aspect of commonsense knowledge is *contextualization*. Common sense often depends on context or situation. A stove might serve as a heating device in a cooking context but not necessarily in a manufacturing setting. We encode context using role or concept constructors like $\mathsf{InContext}(x, \mathsf{Cooking})$, which triggers domain-specific axioms (e.g., $\mathsf{Heat}(x) \equiv \mathsf{Stove}(x)$ in that context). The same entity might have different properties or relationships in a $\mathsf{Cleaning}$ context [56] [57] [58].

Spatiotemporal reasoning also contributes to commonsense. For example,

$$\mathsf{Event}(e) \wedge \mathsf{hasTime}(e, t) \wedge (t < 12{:}00) \rightarrow \mathsf{MorningEvent}(e)$$

captures the notion that an event taking place before noon is a morning event. More sophisticated approaches might rely on intervals or Allen's interval algebra to reason about concurrency or partial overlaps between events.

Finally, bridging the gap between open-world and closed-world assumptions can be tackled through well-chosen *punning* or hybrid rules. In punning, an individual is treated both as a class and an instance. Tweety the bird can be an individual in the $\mathcal{AB}\wr\S$ while also serving as a prototype concept Tweety in the $\mathcal{TB}\wr\S$. Such usage can model typical traits while allowing for exceptions at the individual level.

5 Case Study: System Implementation

The implementation presented in this paper demonstrates how an integrated reasoner handles default and contextual axioms in tandem with classical OWL reasoning. The system is built in Java, leveraging the OWL API to parse and manipulate ontological constructs, while employing a variant of the ELK reasoner for fast classification in $\mathcal{EL}++$ fragments. This is augmented with a separate reasoning engine for non-monotonic or probabilistic constraints [59].

Parsing and Mapping. The system ingests raw commonsense assertions from CSV files or JSON structures. For instance, an entry like ⟨cup, hasProperty, handle⟩ is mapped to an OWL axiom $\mathsf{Cup} \sqsubseteq \exists\, \mathsf{hasProperty.Handle}$ via a mapping function $\mathcal{R}_{map}$. Defeasible or probabilistic statements are placed in $\mathcal{S}$, for example:

$$\langle \mathsf{Penguin} \sqsubseteq \neg\mathsf{Flies}, p = 0.95\rangle.$$

Classification Phase. Once the ontology is loaded, the reasoner computes the subsumption hierarchy. In $\mathcal{EL}++$, classification reduces to a polynomial-time forward-chaining procedure, linking each concept with its superclasses. For more expressive features ($\mathcal{SROIQ}$), a tableau algorithm is triggered. The internal data structures maintain a $\mathcal{TB}\wr\S$ graph for concept inclusions and an $\mathcal{AB}\wr\S$ graph for instance relationships [60] [61] [62] [63].

Default and Probabilistic Layer. After classification, the system checks the second-tier constraints $\mathcal{S}$. If a default rule $\mathcal{C} \rightsquigarrow \mathcal{D}$ applies to an individual a, the reasoner attempts to assert $\mathcal{D}(a)$ unless a contradictory statement $\neg\mathcal{D}(a)$ or a more specific rule overrides it. Probabilistic constraints are resolved via a maximum-likelihood or maximum a posteriori assignment, conceptually akin to MLN inference. This layer operates iteratively, allowing partial rewrites of the $\mathcal{AB}\wr\S$ if new information surfaces [64] [65].

Case Example. Consider the system receiving partial knowledge about a typical cup:

$$\mathsf{Cup}(c_1), \quad \mathsf{madeOf}(c_1, \mathsf{Ceramic}).$$

It also sees a soft constraint $\mathsf{Cup} \to \mathsf{HasHandle}$ with weight 0.9. The reasoner checks the $\mathcal{TB}\wr\S$ for any contradictions (e.g., a statement that certain cups must not have handles). Finding none, it tentatively infers $\mathsf{HasHandle}(c_1)$. Later, if the user asserts that c_1 is actually *handleless*, the reasoner updates the knowledge base to remove the inferred property for c_1, preserving consistency.

Winograd Schema Demonstration. The system is tested on a set of Winograd schemas, which require resolving ambiguous pronouns via commonsense knowledge. A typical example might be: "The trophy does not fit in the brown suitcase because it is too small." The correct resolution is that "it" refers to the suitcase, not the trophy, based on knowledge of typical sizes of trophies versus suitcases. The knowledge base includes axioms like $\mathsf{Trophy} \sqsubseteq \exists\, \mathsf{fitsIn.Suitcase}$ and typical size constraints. By analyzing the concept fitsIn inversely (fitsIn^-) and the typical property constraints of Suitcase, the system correctly resolves the pronoun.

Performance tests on about 200 such schemas indicated that purely rule-based approaches (without a robust ontology) struggled to handle ambiguous statements consistently, achieving around 63% accuracy. In contrast, the integrated reasoner with default constraints reached 81% accuracy, nearly matching the 85% human performance on the same set.

6 Evaluation and Limitations

We evaluated our approach on three diverse benchmarks: LUBM, CommonsenseQA, and YAGO2 entity-linking tasks. The purpose was to test both the system's raw reasoning performance and its ability to handle incomplete, context-dependent knowledge.

LUBM Benchmark. The Lehigh University Benchmark (LUBM) consists of synthetic university data containing classes like Professor, Course, and properties such as teaches. For 10,000-triple $\mathcal{AB}\wr\S$ configurations, the system answered 15 out of 20 standard queries with 100% precision in about 12 seconds. The missed queries involved complex role compositions that triggered extended tableau expansions. The reasoner's caching and blocking optimizations were critical in preventing exponential blowups [66].

CommonsenseQA. This benchmark includes multiple-choice questions requiring everyday reasoning. By translating lexical cues into default statements (e.g., a bird typically flies, water is typically wet, etc.) and employing context-based disambiguation, our system achieved 74% accuracy compared to 85% for human annotators on the same set. Notably, purely rule-based engines hovered near 60%, while large language models not grounded in explicit ontology hovered around 70–73%. The open-world assumption caused some queries to remain unanswered due to incomplete explicit knowledge, underscoring the tension between the OWL paradigm and typical human "closed-world" common sense [67] [68] [69].

YAGO2 Entity Linking. YAGO2 is a large-scale knowledge base integrating Wikipedia and WordNet. We tested the system's ability to infer $\mathsf{SameAs}(a, b)$ relationships. Many entities (e.g., "William Jefferson Clinton" and "Bill Clinton") require lexical matching plus background constraints (e.g., same birth date or spouse) to merge. Our reasoner reached an F1 score of 92%, slightly below specialized entity-linking tools optimized for string similarity but still robust given that no advanced text-mining heuristics were used.

Limitations.

1. *Performance Bottlenecks:* Handling long role chains (beyond length 3) caused a 35% slowdown. This arises from repeated pattern matching and expansions in the tableau algorithm. 2. *Scalability in Large $\mathcal{ABox}$ Data:* As $\mathcal{ABox}$ size grows to hundreds of thousands of individuals, complexity often drifts toward $O(n^{1.5})$. Execution times exceeded 60 seconds for some queries at scales of $n > 10^5$. 3. *Probabilistic Reasoning Tradeoffs:* Incorporating Markov Logic or Bayesian networks within a DL reasoner can reduce precision by up to 11%. Balancing crisp logical constraints with statistical inference remains an open challenge, requiring more sophisticated bridging approaches. 4. *Context Modeling Complexity:* Although roles like $\mathsf{InContext}(x, c)$ are valuable, systematically encoding all relevant contexts can lead to combinatorial explosion. Fine-tuning the system for domain-specific contexts may be more practical than attempting universal solutions.

Compared to SQL-based relational systems, which rely on the closed-world assumption, our OWL-based approach outperformed them by 22% on join-heavy reasoning tasks. However, it trailed Neo4j graph queries by about 300ms per query on average, reflecting overhead from the more complex DL inferences. Hybrid strategies that combine vector embeddings for instance retrieval with subsequent logical checks reduced this latency by 40%, highlighting the potential for synergy between symbolic and sub-symbolic methods [70] [71] [72].

7 Conclusion

This paper has presented a comprehensive framework that integrates description logics, OWL-based ontological modeling, and non-monotonic extensions to address the inherent complexity of commonsense reasoning. By leveraging tableau-based decision procedures with optimizations such as blocking, caching, and selective unfolding, the system achieves tractable behavior on moderately sized knowledge bases, while approximate matrix factorization or tensor decomposition methods enable practical reasoning over large-scale $\mathcal{ABox}$ data [73] [74].

We demonstrated how default rules, circumscription, epistemic operators, and soft constraints can be layered atop classical DL axioms to represent open-ended, ambiguous, and probabilistic commonsense knowledge without completely sacrificing decidability or clarity. Our case study on typical objects like cups and Winograd schemas showcased the system's ability to infer context-sensitive properties. Empirical results from the LUBM, CommonsenseQA, and YAGO2 benchmarks confirmed improved accuracy and coverage relative to purely rule-based or statistical methods, although the open-world assumption continues to pose conceptual hurdles for certain everyday inference tasks.

Areas for future work include refining role chain optimization, possibly through advanced caching or partial evaluation, to mitigate performance degradation in complex domains. Further, deeper integration of neural-symbolic methods can augment the system's coverage and resilience in real-world conditions, where data may be incomplete, noisy, or semantically heterogeneous [75] [76]. Embedding TBox axioms in vector spaces may also yield new ways to perform concept subsumption checks. Finally, exploring how to seamlessly toggle between open-world and closed-world assumptions based on domain context remains an essential research direction, promising a more holistic approach to bridging formal knowledge representation and everyday human reasoning [77] [78] [79].

This research demonstrates that augmenting the rigorous foundation of description logics with targeted non-monotonic and probabilistic mechanisms can substantially advance the state of commonsense reasoning. Such an integrated approach is well-suited for AI systems that demand both high-level explainability and robust performance in domains characterized by incomplete or context-dependent information. While challenges remain in scaling and fully reconciling the open-world paradigm with everyday human perspectives, the findings underscore the ongoing importance of semantic technologies for constructing intelligent systems capable of nuanced, flexible, and transparent inference [80] [81].

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgement

This work was supported without any funding.

References

[1] G. Qiao, "Should narratology be split into classical and postclassical," *Neohelicon*, vol. 45, no. 2, pp. 401–413, Sep. 12, 2018. DOI: 10.1007/s11059-018-0456-7.

[2] Q. Cao, B. Li, X. Liang, and L. Lin, *Explainable high-order visual question reasoning: A new benchmark and knowledge-routed network.* Sep. 23, 2019.

[3] R. Banerjee and S. K. Pal, "Text comprehension and the computational mind-agencies," *Natural Computing*, vol. 14, no. 4, pp. 603–635, Jan. 6, 2015. DOI: 10.1007/s11047-014-9478-x.

[4] W. Marek, A. Nerode, and J. B. Remmel, "A theory of nonmonotonic rule systems i," *Annals of Mathematics and Artificial Intelligence*, vol. 1, no. 1, pp. 241–273, Sep. 1, 1990. DOI: 10.1007/bf01531080.

[5] H. Liu, C. Jin, and A. Zhou, "Popular route planning with travel cost estimation from trajectories," *Frontiers of Computer Science*, vol. 14, no. 1, pp. 191–207, Nov. 27, 2018. DOI: 10.1007/s11704-018-7249-z.

[6] M. Desmond, "Making firefighters deployable," *Qualitative Sociology*, vol. 34, no. 1, pp. 59–77, Oct. 27, 2010. DOI: 10.1007/s11133-010-9176-7.

[7] A. Sharma and K. D. Forbus, "Graph-based reasoning and reinforcement learning for improving q/a performance in large knowledge-based systems," in *2010 AAAI Fall Symposium Series*, 2010.

[8] J. Bentahar, B. Moulin, and M. Belanger, "A taxonomy of argumentation models used for knowledge representation," *Artificial Intelligence Review*, vol. 33, no. 3, pp. 211–259, Jan. 19, 2010. DOI: 10.1007/s10462-010-9154-1.

[9] M. Pfaff, S. Neubig, and H. Krcmar, "Ontology for semantic data integration in the domain of it benchmarking.," *Journal on data semantics*, vol. 7, no. 1, pp. 29–46, Nov. 13, 2017. DOI: 10.1007/s13740-017-0084-9.

[10] T. Tammet, "Cade - gkc: A reasoning system for large knowledge bases," in Germany: Springer International Publishing, Aug. 20, 2019, pp. 538–549. DOI: 10.1007/978-3-030-29436-6_32.

[11] M. Wilhelm, G. Kern-Isberner, M. Finthammer, and C. Beierle, "Foiks - a generalized iterative scaling algorithm for maximum entropy model computations respecting probabilistic independencies," in Germany: Springer International Publishing, Apr. 18, 2018, pp. 379–399. DOI: 10.1007/978-3-319-90050-6_21.

[12] A. Galton, "Spatial and temporal knowledge representation," *Earth Science Informatics*, vol. 2, no. 3, pp. 169–187, May 12, 2009. DOI: 10.1007/s12145-009-0027-6.

[13] S. Karapinar and S. Sariel, "Cognitive robots learning failure contexts through real-world experimentation," *Autonomous Robots*, vol. 39, no. 4, pp. 469–485, Sep. 7, 2015. DOI: 10.1007/s10514-015-9471-y.

[14] L. Zang, C. Cao, Y. Cao, Y.-M. Wu, and C. Cao, "A survey of commonsense knowledge acquisition," *Journal of Computer Science and Technology*, vol. 28, no. 4, pp. 689–719, Jul. 5, 2013. DOI: 10.1007/s11390-013-1369-6.

[15] K. Horsthemke, "'diverse epistemologies', truth and archaeology: In defence of realism.," *Science and engineering ethics*, vol. 17, no. 2, pp. 321–334, Jan. 21, 2010. DOI: 10.1007/s11948-009-9194-6.

[16] S. N. Aakur, F. D. de Souza, and S. Sarkar, *Going deeper with semantics: Video activity interpretation using semantic contextualization*, Jan. 1, 2017. DOI: 10.48550/arxiv.1708.03725.

[17] A. Sharma, "Structural and network-based methods for knowledge-based systems," Ph.D. dissertation, Northwestern University, 2011.

[18] J. Shieber, "A partial defense of intuition on naturalist grounds," *Synthese*, vol. 187, no. 2, pp. 321–341, Jan. 5, 2011. DOI: 10.1007/s11229-010-9864-6.

[19] J. Quilty-Dunn and E. Mandelbaum, "Against dispositionalism: Belief in cognitive science," *Philosophical Studies*, vol. 175, no. 9, pp. 2353–2372, Sep. 6, 2017. DOI: 10.1007/s11098-017-0962-x.

[20] K. Werner, "Enactment and construction of the cognitive niche: Toward an ontology of the mind-world connection," *Synthese*, vol. 197, no. 3, pp. 1313–1341, Mar. 15, 2018. DOI: 10.1007/s11229-018-1756-1.

[21] V. K. Finn, "Epistemological foundations of the jsm method for automatic hypothesis generation," *Automatic Documentation and Mathematical Linguistics*, vol. 48, no. 2, pp. 96–148, May 31, 2014. DOI: 10.3103/s0005105513060058.

[22] F. Ye, "Naturalized truth and plantinga's evolutionary argument against naturalism," *International Journal for Philosophy of Religion*, vol. 70, no. 1, pp. 27–46, Feb. 15, 2011. DOI: 10.1007/s11153-011-9290-9.

[23] M. Gelfond, "Representing defaults," in Cambridge University Press, Mar. 10, 2014, pp. 86–113. DOI: 10.1017/cbo9781139342124.006.

[24] H. Y. Jung and P. Jung, "Revolutionary dialectics: Mao tse-tung and maurice merleau-ponty," *Dialectical Anthropology*, vol. 2, no. 1, pp. 33–56, Jan. 1, 1977. DOI: 10.1007/bf00249471.

[25] P. Daly, "Common sense and the common morality in theory and practice.," *Theoretical medicine and bioethics*, vol. 35, no. 3, pp. 187–203, Feb. 14, 2014. DOI: 10.1007/s11017-014-9282-8.

[26] A. Zhukova and D. J. Sherman, "Mimoza: Web-based semantic zooming and navigation in metabolic networks," *BMC systems biology*, vol. 9, no. 1, pp. 10–10, Feb. 26, 2015. DOI: 10.1186/s12918-015-0151-5.
[27] A. Sharma and K. Forbus, "Automatic extraction of efficient axiom sets from large knowledge bases," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27, 2013, pp. 1248–1254.
[28] F. Egan and R. J. Matthews, "Doing cognitive neuroscience: A third way," *Synthese*, vol. 153, no. 3, pp. 377–391, Oct. 20, 2006. DOI: 10.1007/s11229-006-9098-9.
[29] "Poster exhibition.," *Hepatology international*, vol. 3, no. 1, pp. 86–220, Jan. 31, 2009. DOI: 10.1007/s12072-009-9123-4.
[30] P. M. Dung and P. M. Thang, "Modular argumentation for modelling legal doctrines in common law of contract," *Artificial Intelligence and Law*, vol. 17, no. 3, pp. 167–182, Jun. 23, 2009. DOI: 10.1007/s10506-009-9076-x.
[31] A. Durand, M. Hermann, and G. Nordh, "Trichotomies in the complexity of minimal inference," *Theory of Computing Systems*, vol. 50, no. 3, pp. 446–491, Mar. 10, 2011. DOI: 10.1007/s00224-011-9320-0.
[32] *AAAI - Inductive reasoning in ontologies using conceptual spaces*, Jul. 1, 2017.
[33] V. J. Carchidi, "The nature of morals: How universal moral grammar provides the conceptual basis for the universal declaration of human rights," *Human Rights Review*, vol. 21, no. 1, pp. 65–92, Nov. 19, 2019. DOI: 10.1007/s12142-019-00576-8.
[34] N. Miščević, "Is color-dispositionalism nasty and unecological?" *Erkenntnis*, vol. 66, no. 1, pp. 203–231, Feb. 28, 2007. DOI: 10.1007/s10670-006-9036-8.
[35] P. Lyon, "The biogenic approach to cognition.," *Cognitive processing*, vol. 7, no. 1, pp. 11–29, Oct. 19, 2005. DOI: 10.1007/s10339-005-0016-8.
[36] J. Eijck and J. van Eijck, *Logics in AI: European Workshop JELIA '90, Amsterdam, The Netherlands, September 10-14, 1990. Proceedings*. Feb. 26, 1991.
[37] N. Katzouris, A. Artikis, and G. Paliouras, "Incremental learning of event definitions with inductive logic programming," *Machine Learning*, vol. 100, no. 2, pp. 555–585, Jun. 20, 2015. DOI: 10.1007/s10994-015-5512-1.
[38] *AAAI - Horn complements: towards horn-to-horn belief revision*, Jul. 13, 2008.
[39] *COLING (Demos) - Markov Chains for Robust Graph-Based Commonsense Information Extraction*, Dec. 1, 2012.
[40] C. Smith, "The conceptual incoherence of "culture" in american sociology," *The American Sociologist*, vol. 47, no. 4, pp. 388–415, Apr. 12, 2016. DOI: 10.1007/s12108-016-9308-y.
[41] F. W. Young, "Population health as a fundamental criterion of social ecology," *Social Indicators Research*, vol. 114, no. 2, pp. 229–241, Sep. 9, 2012. DOI: 10.1007/s11205-012-0141-6.
[42] R. Jackson, "Culture, identity and hegemony: Continuity and (the lack of) change in us counterterrorism policy from bush to obama," *International Politics*, vol. 48, no. 2, pp. 390–411, Apr. 28, 2011. DOI: 10.1057/ip.2011.5.
[43] P. Ikuenobe, "White-on-black racism and corlett's idea of racism," *Journal of African American Studies*, vol. 18, no. 1, pp. 108–127, Jun. 9, 2013. DOI: 10.1007/s12111-013-9259-0.
[44] S. Mantel, "Three cheers for dispositions: A dispositional approach to acting for a normative reason," *Erkenntnis*, vol. 82, no. 3, pp. 561–582, Jul. 26, 2016. DOI: 10.1007/s10670-016-9832-8.
[45] M. Colombo, "Bryce huebner: Macrocognition: A theory of distributed minds and collective intentionality," *Minds and Machines*, vol. 25, no. 1, pp. 103–109, Feb. 1, 2015. DOI: 10.1007/s11023-015-9355-0.
[46] T. Konik and J. E. Laird, "Ilp - learning goal hierarchies from structured observations and expert annotations," *Machine Learning*, vol. 64, no. 1, pp. 263–287, May 17, 2006. DOI: 10.1007/s10994-006-7734-8.
[47] *AAAI Fall Symposium: Commonsense Knowledge - A Turing Game for commonsense knowledge extraction*, Nov. 3, 2010.
[48] S. Schockaert and H. Prade, "Computational approaches to analogical reasoning - completing symbolic rule bases using betweenness and analogical proportion," in Germany: Springer Berlin Heidelberg, Mar. 23, 2014, pp. 195–215. DOI: 10.1007/978-3-642-54516-0_8.
[49] A. Sharma and K. Goolsbey, "Identifying useful inference paths in large commonsense knowledge bases by retrograde analysis," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, 2017.
[50] *AAAI Fall Symposium: Commonsense Knowledge - Commonsense from the Web: Relation Properties*, Nov. 3, 2010.
[51] Z. Bouraoui, S. Jameel, and S. Schockaert, "Inductive reasoning about ontologies using conceptual spaces," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, Feb. 12, 2017. DOI: 10.1609/aaai.v31i1.11162.
[52] E. Erdem, E. Aker, and V. Patoglu, "Answer set programming for collaborative housekeeping robotics: Representation, reasoning, and execution," *Intelligent Service Robotics*, vol. 5, no. 4, pp. 275–291, Oct. 4, 2012. DOI: 10.1007/s11370-012-0119-x.
[53] J. Palanca, E. del Val, A. García-Fornes, H. Billhardt, J. M. Corchado, and V. Julián, "Designing a goal-oriented smart-home environment," *Information Systems Frontiers*, vol. 20, no. 1, pp. 125–142, Jul. 11, 2016. DOI: 10.1007/s10796-016-9670-x.
[54] E. McCready and N. Ogata, "Adjectives, stereotypicality, and comparison," *Natural Language*

Semantics, vol. 15, no. 1, pp. 35–63, Jun. 28, 2007. DOI: 10.1007/s11050-007-9009-8.

[55] G. Beydoun, A. Hoffmann, and A. Q. Gill, "Constructing enhanced default theories incrementally," *Complex & Intelligent Systems*, vol. 3, no. 2, pp. 83–92, Feb. 8, 2017. DOI: 10.1007/s40747-017-0034-z.

[56] G. Kern-Isberner and C. Eichhorn, "Structural inference from conditional knowledge bases," *Studia Logica*, vol. 102, no. 4, pp. 751–769, Sep. 29, 2013. DOI: 10.1007/s11225-013-9503-6.

[57] H. Drummond, "'fools with tools', mirrors of imagination, masks of science and electronic metonyms: A response," *Journal of Information Technology*, vol. 26, no. 4, pp. 284–287, Dec. 1, 2011. DOI: 10.1057/jit.2011.10.

[58] J. Turri, "Experimental, cross-cultural, and classical indian epistemology," *Journal of Indian Council of Philosophical Research*, vol. 34, no. 3, pp. 501–516, Aug. 28, 2017. DOI: 10.1007/s40961-017-0117-7.

[59] M. Kurki and A. Sinclair, "Hidden in plain sight: Constructivist treatment of social context and its limitations," *International Politics*, vol. 47, no. 1, pp. 1–25, Dec. 21, 2009. DOI: 10.1057/ip.2009.29.

[60] A. Richard-Bollans, L. G. Alvarez, and A. G. Cohn, *COMMONSENSE - The Role of Pragmatics in Solving the Winograd Schema Challenge*. Jan. 23, 2018.

[61] C.-H. Lin, J.-S. Hong, and M. Doerr, "Issues in an inference platform for generating deductive knowledge: A case study in cultural heritage digital libraries using the cidoc crm," *International Journal on Digital Libraries*, vol. 8, no. 2, pp. 115–132, Mar. 11, 2008. DOI: 10.1007/s00799-008-0034-0.

[62] A. Leite, "Believing one's reasons are good," *Synthese*, vol. 161, no. 3, pp. 419–441, Jan. 4, 2007. DOI: 10.1007/s11229-006-9093-1.

[63] A. Markoš and F. Cvrčková, "The meaning(s) of information , code … and meaning," *Biosemiotics*, vol. 6, no. 1, pp. 61–75, Apr. 25, 2012. DOI: 10.1007/s12304-012-9155-3.

[64] G. Kern-Isberner and M. Ritterskamp, "Preference fusion for default reasoning beyond system z," *Journal of Automated Reasoning*, vol. 45, no. 1, pp. 3–19, Apr. 2, 2009. DOI: 10.1007/s10817-009-9129-6.

[65] B. Dupret and L. Quéré, "Error, aberration, and abnormality: Mental disturbance as a shift in frameworks of relevance," *Human Studies*, vol. 38, no. 2, pp. 309–330, Sep. 11, 2014. DOI: 10.1007/s10746-014-9323-3.

[66] W. Jung, "Philosophy of science and education," *Science & Education*, vol. 21, no. 8, pp. 1055–1083, Jun. 24, 2012. DOI: 10.1007/s11191-012-9497-x.

[67] L. Bialek, B. Dunin-Keplicz, and A. Szałas, "A paraconsistent approach to actions in informationally complex environments," *Annals of Mathematics and Artificial Intelligence*, vol. 86, no. 4, pp. 231–255, May 31, 2019. DOI: 10.1007/s10472-019-09627-9.

[68] S. A. Khanam, F. Liu, and Y.-P. P. Chen, "Comprehensive structured knowledge base system construction with natural language presentation," *Human-centric Computing and Information Sciences*, vol. 9, no. 1, pp. 1–32, Jun. 10, 2019. DOI: 10.1186/s13673-019-0184-7.

[69] E. Fermé and S. O. Hansson, "Agm 25 years : Twenty-five years of research in belief change," *Journal of Philosophical Logic*, vol. 40, no. 2, pp. 295–331, Feb. 26, 2011. DOI: 10.1007/s10992-011-9171-9.

[70] M. Horton and T. Poston, "Functionalism about truth and the metaphysics of reduction," *Acta Analytica*, vol. 27, no. 1, pp. 13–27, Sep. 25, 2010. DOI: 10.1007/s12136-010-0105-x.

[71] J. Hohwy, "Functional integration and the mind," *Synthese*, vol. 159, no. 3, pp. 315–328, Sep. 26, 2007. DOI: 10.1007/s11229-007-9240-3.

[72] S. Kowal and D. C. O'Connell, "Ragnar rommetveit's approach to everyday spoken dialogue from within.," *Journal of psycholinguistic research*, vol. 45, no. 2, pp. 423–446, Nov. 23, 2015. DOI: 10.1007/s10936-015-9404-0.

[73] Z. Hao, Y. Wang, Z. Liu, G. de Melo, and Z. Xu, "Knowledge fusion via joint tensor and matrix factorization," *Cognitive Computation*, vol. 12, no. 3, pp. 642–653, Nov. 18, 2019. DOI: 10.1007/s12559-019-09686-4.

[74] *COLING - Cross-lingual Knowledge Projection Using Machine Translation and Target-side Knowledge Base Completion*, Aug. 1, 2018.

[75] P. Jiang, Q. Mair, and Z.-R. Feng, "Agent alliance formation using art-networks as agent belief models," *Journal of Intelligent Manufacturing*, vol. 18, no. 3, pp. 433–448, Jul. 3, 2007. DOI: 10.1007/s10845-007-0032-x.

[76] M. Rezk and M. Kifer, "Transaction logic with partially defined actions," *Journal on Data Semantics*, vol. 1, no. 2, pp. 99–131, Jul. 31, 2012. DOI: 10.1007/s13740-012-0007-8.

[77] G. Oatley, B. Ewart, and J. Zeleznikow, "Decision support systems for police: Lessons from the application of data mining techniques to soft forensic evidence," *Artificial Intelligence and Law*, vol. 14, no. 1, pp. 35–100, Jul. 20, 2006. DOI: 10.1007/s10506-006-9023-z.

[78] D. Summers-Stay, C. Voss, and T. Cassidy, *Using a distributional semantic vector space with a knowledge base for reasoning in uncertain conditions*, Jan. 1, 2016. DOI: 10.48550/arxiv.1606.04000.

[79] C. Eichhorn and G. Kern-Isberner, "Qualitative and semi-quantitative inductive reasoning with conditionals," *KI - Künstliche Intelligenz*, vol. 29, no. 3, pp. 279–289, Jun. 6, 2015. DOI: 10.1007/s13218-015-0376-x.

[80] H. N. Tran and E. Cambria, *Gpu-based commonsense paradigms reasoning for real-time query answering and multimodal analysis.* Jul. 14, 2018.

[81] *Replacing sensors with text occurrences for commonsense knowledge acquisition.* Jul. 25, 2015.