Morphpublishing
J. Comput. Intell. Mach. Reason. Decis.-Mak.
open access academic journal
Peer reviewed

Comparative Evaluation of Graph, Hierarchical, and Relational Structuring Paradigms in Large Commonsense Knowledge Bases

Omar Alfarouq[a], Ahmad Alnaser[b]

Abstract: Commonsense knowledge bases (CSKBs) serve as foundational infrastructures for enabling machines to reason about everyday concepts. While numerous structuring paradigms exist, their comparative efficacy in balancing expressivity, computational efficiency, and scalability remains underexplored. This paper presents a systematic evaluation of graph-based, hierarchical, and relational database approaches for organizing large-scale CSKBs. Graph representations model knowledge as nodes and edges, enabling flexible traversal but incurring overhead in path-sensitive queries. Hierarchical methods, such as taxonomies, optimize inheritance reasoning but struggle with cross-domain ambiguity. Relational paradigms, grounded in formal set theory, enforce strict normalization constraints that enhance integrity but limit dynamic expansion. We formalize each paradigm using algebraic structures, predicate logic, and complexity-theoretic metrics. A novel hypergraph-based hybrid model is proposed to mitigate rigidity in relational systems while preserving hierarchical inheritance. Experiments on ConceptNet, WordNet, and a proprietary dataset quantify throughput for subsumption, adjacency, and transitivity queries. Results indicate relational systems outperform others in conjunctive queries (F1: 0.92 vs. 0.78 for graphs) but exhibit exponential latency growth during schema revisions. Graph-based approaches achieve high recall in path queries but require heavier indexing. Hierarchical systems optimize memory usage but suffer precision loss in multi-inheritance contexts. The hybrid model reduces transitivity errors by 37% via constraint-driven edge weighting. This work offers framework-agnostic optimization guidelines for CSKB engineers.

1. Introduction

Humans rely on commonsense knowledge to navigate routine tasks and interpret contextual cues. As artificial intelligence systems increasingly integrate into complex real-world environments, they require a robust layer of commonsense reasoning to disambiguate user queries, infer implicit relationships, and adapt to novel situations.

[a]Al al-Bayt University, Department of Computer Science, Mafraq Highway, Mafraq, 25113, Jordan.
[b]Mutah University, Faculty of Information Technology, Queen Rania Street, Al-Karak, 61710, Jordan.

Commonsense knowledge bases (CSKBs) encode such everyday facts, ranging from physical interactions (e.g., "bottles contain liquids") to sociocultural norms (e.g., "shaking hands is a greeting"). Although large-scale CSKBs like ConceptNet, WordNet, and YAGO have existed for years, the internal representation structure used to store and retrieve their knowledge can drastically affect performance, scalability, and flexibility [1] [2].

A fundamental challenge in CSKB construction is the trade-off between structured formalism and expressive flexibility. Graph-based representations, such as those used in ConceptNet, store commonsense facts as triples (subject, relation, object), facilitating efficient reasoning and retrieval. These representations excel in relational inference but often struggle with capturing nuanced contextual variations. On the other hand, distributional representations, including embeddings trained on large text corpora, model commonsense as high-dimensional vector spaces where semantic similarities emerge through proximity. While these embeddings provide adaptability, they inherently lack explicit, human-interpretable relations [3] [4] [5, 6]. A hybrid approach that fuses structured graphs with distributional semantics can enhance both interpretability and generalizability. Recent advances leverage transformer-based language models to dynamically retrieve and synthesize commonsense knowledge, bridging the gap between static knowledge bases and adaptive contextual understanding.

The acquisition of commonsense knowledge poses additional difficulties due to its implicit nature. Unlike encyclopedic facts, which can be extracted from structured databases or textual corpora, commonsense knowledge is often unstated because it is assumed to be universally known. As a result, automated methods must infer such knowledge from indirect signals. Crowdsourcing has played a pivotal role in building CSKBs, with projects like ATOMIC relying on human annotators to contribute inferential commonsense relationships (e.g., "If someone apologizes, they likely feel regret"). However, this approach is labor-intensive and susceptible to biases. Alternatively, self-supervised learning methods have emerged, where language models distill commonsense knowledge by predicting missing information in large corpora. These methods reduce human effort but introduce challenges in filtering hallucinated or nonfactual inferences.

Evaluating the quality of commonsense knowledge representations requires robust benchmarks. Metrics such as coverage, coherence, and applicability help assess the utility of a CSKB. Coverage quantifies the breadth of knowledge encoded within a resource, ensuring that diverse domains such as physics, psychology, and social conventions are represented. Coherence measures the logical consistency of stored facts, identifying contradictions or redundancies that could impair reasoning. Applicability evaluates how effectively a CSKB supports downstream tasks, including question answering, dialogue generation, and robotic planning. Standardized datasets like CSQA (CommonsenseQA) and PIQA (Physical Interaction QA) serve as benchmarks for assessing the practical impact of commonsense reasoning models [7] [8] [9].

A comparison of major CSKBs in terms of structure, coverage, and applications is shown in Table 1.

Advancements in neural symbolic reasoning have further enhanced commonsense knowledge integration. Symbolic AI, which traditionally relies on logic-based representations, has been combined with neural architectures to improve inference capabilities. Neuro-symbolic systems allow explicit rule-based commonsense reasoning while leveraging deep learning for generalization. This hybrid approach improves the adaptability of AI systems in novel scenarios, where purely symbolic methods lack flexibility, and purely neural methods struggle with interpretability [10] [11] [12] [13].

Real-world deployment of commonsense reasoning systems faces practical obstacles, including computational efficiency and domain-specific adaptation. Large-scale knowledge bases require optimized storage and retrieval mechanisms to ensure real-time applicability. Indexing techniques such as graph embeddings, sparse attention mechanisms, and knowledge distillation have been employed to enhance efficiency. Additionally, domain adaptation

Table 1. Comparison of Major Commonsense Knowledge Bases

CSKB	Representation	Coverage	Application Domains
ConceptNet	Graph-based triples	Broad (linguistic, social, physical)	NLP, chatbots, reasoning tasks
WordNet	Lexical hierarchy	Linguistic focus	Word sense disambiguation, ontology construction
YAGO	Ontological knowledge graph	Encyclopedic and factual	Semantic search, linked data
ATOMIC	Event-based inferential triples	Social and psychological norms	Narrative understanding, human-AI interaction

strategies, including transfer learning and few-shot learning, allow AI systems to specialize commonsense reasoning for medical, legal, or scientific applications. Context-aware commonsense integration remains a crucial research direction, where AI systems dynamically adjust their reasoning strategies based on situational cues.

A deeper analysis of commonsense reasoning integration in AI applications is presented in Table 2.

Table 2. Commonsense Knowledge in AI Applications

Application Domain	Commonsense Requirements	Examples
Conversational AI	Understanding implicit intents, contextual inference	Chatbots, virtual assistants
Robotics	Physical world interaction, causal reasoning	Autonomous navigation, human-robot collaboration
Healthcare AI	Medical commonsense, patient interaction reasoning	Clinical decision support, symptom checking
Autonomous Systems	Scenario prediction, hazard identification	Self-driving cars, risk assessment

There are three widely adopted paradigms for structuring large CSKBs. The first involves graph-based architectures, which typically represent concepts as vertices, relationships as edges, and sometimes incorporate weights or labels for fine-grained semantics. These methods are amenable to multi-hop traversals using adjacency matrices or lists, and they can capture irregular structures well. However, large graphs can introduce complex cycles and high storage overhead when indexing adjacency lists or matrices. A second paradigm arranges concepts in hierarchical or taxonomic structures, imposing partial orders to accelerate subsumption queries. Hierarchical methods provide efficient single-inheritance relationships and can exploit specialized indexing to answer questions about superclass and subclass relations. Yet they tend to suffer under multi-inheritance or cross-domain expansions, requiring ad-hoc workarounds that can degrade precision [14] [15].

Relational approaches form the third paradigm, modeling concepts and relations in tables under strict normalization principles. By enforcing formal schemas, they provide robust integrity checks, concurrency control, and well-studied query optimization strategies. Nonetheless, relational designs can be cumbersome for capturing highly dynamic or semistructured commonsense facts. Multi-hop reasoning often requires multiple self-joins and can degrade performance if the database schema is not carefully tuned [16] [17] [18] [19, 20].

Selecting among these paradigms is nontrivial. Each provides trade-offs in logic expressivity, update efficiency, memory usage, and typical query latency. For instance, a graph-based CSKB might excel at answering adjacency or path queries but becomes unwieldy if the system must frequently perform deeper classification or enforce complex constraints. A hierarchical approach is elegant for consistent taxonomies with well-defined inheritance paths, but handling cyclical or cross-branch relationships can demand significant engineering overhead. Relational databases are highly mature from a software engineering perspective, with powerful transactional guarantees, yet they can prove restrictive in representing the open-ended nature of commonsense facts [21] [22] [23]. To further illustrate the strengths and weaknesses of these paradigms, Table 3 provides a comparative analysis based on key operational and structural characteristics.

Table 3. Comparison of CSKB Structuring Paradigms

Paradigm	**Strengths**	**Weaknesses**	**Best Use Cases**
Graph-based	Flexible structure, efficient for adjacency queries, supports multi-hop reasoning	High storage overhead, complex indexing, potential cyclic dependencies	Commonsense inference, knowledge discovery
Hierarchical	Efficient single-inheritance, fast subsumption queries, structured organization	Poor handling of multi-inheritance, limited cross-domain adaptability	Taxonomies, structured domain ontologies
Relational	Strong integrity constraints, optimized querying, concurrency control	Rigid schemas, difficult to model semi-structured facts, performance bottlenecks in deep reasoning	Well-defined, static domains requiring data consistency

Given these trade-offs, recent developments in hybrid CSKB architectures have sought to leverage the advantages of multiple paradigms. For example, hybrid models often store structured relational data alongside graph-based indices, allowing flexible traversal without sacrificing query optimization. Additionally, embeddings derived from deep learning models have been integrated into CSKBs to provide a latent semantic representation of commonsense knowledge. This enables AI systems to interpolate between explicit relational structures and implicit knowledge captured from unstructured text corpora [24] [25] [26].

The evolution of commonsense reasoning frameworks has also led to the development of modular architectures where different components specialize in specific reasoning tasks. A knowledge graph module may handle connectivity and entity resolution, while a hierarchical layer provides structured taxonomy-based inferences, and a neural embedding module enhances flexibility by capturing contextual variations. These modular approaches offer a promising direction for scalable CSKBs, allowing fine-tuned performance adjustments depending on the application requirements.

Despite these advancements, ensuring robustness in commonsense knowledge integration remains a challenge. Real-world AI applications often require rapid updates to accommodate newly emerging knowledge, and rigid database structures may struggle with the dynamic nature of commonsense facts. Streaming updates and incremental learning strategies are being explored to allow CSKBs to evolve continuously without requiring full retraining or expensive re-indexing. Moreover, explainability remains a key concern, as opaque reasoning processes can hinder user trust and adoption in high-stakes domains like healthcare, law, and autonomous systems [27] [28] [29].

A particularly important consideration in structuring CSKBs is their adaptability to multi-modal reasoning. Traditional knowledge bases primarily focus on text-based representations, but integrating visual, auditory, and sensory data sources enhances their ability to model real-world contexts. Multi-modal commonsense learning seeks to bridge this gap by fusing linguistic knowledge with perceptual signals, thereby improving AI's ability to interpret human interactions, recognize affordances in objects, and reason about cause-effect relationships in the physical world.

Table 4 summarizes various hybrid CSKB strategies, highlighting their architectural composition and typical advantages.

Table 4. Hybrid CSKB Strategies and Their Advantages

Hybrid Model Type	Architectural Composition	Key Advantages
Graph-Relational Hybrid	Graph-based reasoning combined with relational data integrity	Enables flexible traversal while maintaining strict data consistency
Neural-Symbolic Integration	Deep learning embeddings combined with rule-based logic	Enhances adaptability and inference while preserving explicit reasoning
Multi-Modal Commonsense Model	Text-based CSKBs integrated with visual and sensory data	Improves contextual reasoning in real-world scenarios

To elucidate these concerns, this paper examines each paradigm from multiple perspectives. A formal review includes algebraic characterizations of adjacency matrices and transitive closures, hierarchical partial orders, and relational joins. Complexity analyses illustrate the computational cost of common CSKB operations, such as insertion of new concepts, resolution of multi-hop queries, and reorganization of schema definitions. Furthermore, benchmark experiments on real and synthetic data shed light on trade-offs in accuracy, latency, and resource consumption across graph-based, hierarchical, and relational representations [30] [31] [32].

In addition, the paper introduces a hybrid hypergraph-based architecture that blends the strengths of different paradigms to mitigate their individual shortcomings. Specifically, the approach constructs a set of typed hyperedges for more flexible connectivity, coupled with a partial hierarchy for consistent taxonomic reasoning. This model aims to preserve the efficient inheritance features found in strictly hierarchical systems while simultaneously providing the lateral flexibility of graph-based designs and the rigorous constraints found in relational databases.

Section-by-section, the discussion proceeds as follows. First, key concepts of graph-based knowledge structuring are explored, highlighting the role of adjacency matrices, label semantics, and eigenvalue-based computations. The next section delves into hierarchical representations, focusing on partial orders, single- and multi-inheritance constraints, and the management of cross-domain ambiguities. Relational paradigms for CSKBs then come to the forefront, introducing the notion of normal forms, the cost of join operations, and the unique challenges of dynamic schema evolution. An in-depth comparative analysis follows, featuring formal complexity theorems and a conceptual mapping of each paradigm's strengths. Building on these observations, the paper proposes a hybrid model, showing how hypergraph components can be integrated with hierarchical structures for improved multi-hop reasoning and controlled redundancy. Finally, an experimental evaluation is presented, detailing the empirical performance of each paradigm under diverse query workloads, followed by a conclusion summarizing key findings and directions for future research [33] [34] [35, 36].

This introduction has underscored the motivation and objectives underlying the study, emphasizing the importance of representation choices in the design of CSKBs. The subsequent sections develop these themes in detail, providing both theoretical rigor and real-world benchmarks. Readers seeking to deploy or refine large-scale commonsense reasoning systems will find actionable insights, as well as pointers to advanced optimization techniques.

2. Graph-Based Knowledge Structuring

Graph representations conceptualize CSKBs as directed (or undirected) structures composed of nodes and edges. Formally, let a knowledge graph be denoted by $G = (V, E)$, where each vertex $v_i \in V$ corresponds to a concept or entity, and each edge $e_{ij} \in E$ signifies a relationship from v_i to v_j. Edges often carry labels $l(e_{ij}) \in \mathcal{L}$ to denote relation types such as `IsA`, `PartOf`, or `Causes`. Weights w_{ij} in $[0, 1]$ can capture confidence scores or semantic similarity.

A fundamental operation in many graph-based CSKBs is the multi-hop query, which asks whether a path exists from one concept to another through intermediate relations. To handle such queries systematically, adjacency matrices are often used. Define an $|V| \times |V|$ matrix A such that $A_{ij} = w_{ij}$ if there is an edge from v_i to v_j, and $A_{ij} = 0$ otherwise. Multi-hop paths of length k can be tracked by computing A^k, using semiring operations $\oplus$ and $\otimes$ if the domain extends beyond simple Boolean logic. For example, in a probabilistic setting, $\otimes$ might be multiplication for combining the probabilities of consecutive edges, while $\oplus$ might be a max or sum operator for parallel path aggregation.

An illustrative logic statement for adjacency can be written as:

$$(\exists v_k)\big((v_i, v_k) \in E \wedge (v_k, v_j) \in E\big) \quad \Longleftrightarrow \quad (A^2)_{ij} \neq 0.$$

In the case of length-2 paths, the existence of at least one intermediate vertex v_k fulfilling the adjacency condition implies a nonzero entry in the squared adjacency matrix. Extending this logic to higher-order paths becomes computationally more expensive, particularly when cycles exist in the graph. Cycles can inflate the spectral radius $\rho(A)$, leading to slower convergence in iterative algorithms that approximate A^k or attempt to compute transitive closures. Jordan normal forms and eigenvalue decomposition can theoretically accelerate matrix exponentiation, but implementing these techniques at scale is resource-intensive.

From a storage perspective, graph-based systems often adopt adjacency lists instead of dense matrices, especially when the knowledge graph is large but relatively sparse. If $|E|$ is on the order of m, adjacency lists can reduce space complexity relative to the $O(n^2)$ overhead of storing a full matrix. Compressed representations, such as Huffman-coded edge lists, exploit the frequency distribution of edges. However, real-world CSKBs can still accumulate millions of edges, making memory consumption a critical issue. Dynamic updates—such as inserting a new concept v_{new} with d edges—must also be handled. An update may require pointer manipulations in adjacency lists, leading to fragmentation or reallocation overheads that degrade performance over time [37] [38] [39].

Another challenge lies in ensuring consistency and reducing noise in graph-based CSKBs. Because new edges or relationships might be introduced from automated extractors or crowd-sourcing, contradictory facts can appear. Algorithms that prune low-weight edges or attempt to unify semantically redundant vertices can partially address this. Constraints can be expressed in propositional or predicate logic. A sample rule for conflict detection might read:

$$\neg((v_i, v_j) \in E \land l(e_{ij}) = \texttt{IsA} \land l(e_{ij}) = \texttt{NotA}).$$

In words, no vertex pair can simultaneously hold an IsA relation and a NotA relation of the same type. While straightforward, systematically detecting these conflicts at scale can be nontrivial.

Despite these drawbacks, graph-based representations remain popular due to their flexibility. Highly interconnected knowledge domains, such as cause-effect or part-whole relationships spanning multiple semantic categories, benefit from unconstrained adjacency. Graph traversal algorithms, from simple breadth-first searches to more advanced label-propagation and embedding techniques, can glean insights from the global connectivity structure. For instance, random walk–based embeddings or spectral methods that factorize A into lower-dimensional representations can discover latent groupings of concepts for downstream tasks.

In typical benchmarks, graph-based CSKBs handle multi-hop adjacency queries efficiently as long as the number of steps k remains small. Performance degrades with deeper queries because the search space expands exponentially in the number of possible paths. Meanwhile, updates occur relatively smoothly if the concurrency level is not extreme. Systems such as ConceptNet demonstrate the viability of large-scale graph-based commonsense reasoning, offering integration with natural language processing pipelines. However, as we shall see in later sections, hierarchical or relational paradigms may outperform graph-based systems in specific workloads that emphasize rigid classification or frequent relational joins [40] [41] [42].

3. Hierarchical Knowledge Organization

Hierarchical or taxonomic structures impose a partial order on concepts, capturing subsumption (e.g., "an apple is a kind of fruit") and enabling consistent inheritance. Formally, define a hierarchy T as a directed acyclic graph in which each node v_i has at most one immediate parent, enforcing single inheritance. A more general structure may allow multiple parents but at the cost of complicating inheritance logic.

Let $v_i \preceq v_j$ denote that concept v_i is a subclass (or sub-concept) of v_j. In a tree-based CSKB, each node has a path to a root concept, and queries about whether v_i is subsumed by v_j involve tracing upward links until either v_j is encountered or the root is reached. The complexity for such a lookup in a balanced tree is $O(\log n)$. However, many knowledge hierarchies are not strictly balanced, and some subtrees may grow disproportionately.

A logic statement capturing subsumption might read:

$$(v_i \preceq v_j) \land (v_j \preceq v_k) \implies (v_i \preceq v_k).$$

Transitivity is a cornerstone of hierarchical reasoning, ensuring that membership in a lower-level category implies membership in all ancestor categories. In systems like WordNet, these transitive properties speed up "hypernym queries" (queries about superclasses of a concept), which can be resolved via path tracing. To accelerate repeated queries, many implementations store memoized paths or transitive closures so that subsumption checks become constant-time lookups.

However, multi-inheritance arises naturally in domains with overlapping categories. For instance, "tomato" might be simultaneously classified as a vegetable and a fruit, leading to diamonds in the inheritance graph. Formally, for a node v_i with parents v_a and v_b, these parents might share their own ancestors or attributes, and merging them becomes nontrivial. Some taxonomic CSKBs address this by replicating subtrees or introducing intersection nodes $v_a \sqcap v_b$, but such maneuvers can bloat the hierarchy and erode semantic clarity [43] [44] [45].

A typical disadvantage of purely hierarchical representations is limited support for lateral relationships. If a concept has many peer-level connections, the hierarchy cannot capture them directly without adding cross-links, which break the acyclic property. Incorporating these links either transforms the structure into a more general graph or forces the creation of "pseudo-hierarchies." This can blur the conceptual neatness of a single-inheritance tree.

Memory usage in hierarchical CSKBs can be reduced through entropy-aware encodings. Frequency-based compression assigns shorter codes to more common concepts, leveraging Huffman-like schemes to store node identifiers. Such compression is especially useful if the system stores repeated paths or partial subtrees for fast lookups. Symbolically, if $f(v_i)$ denotes the frequency of concept v_i, then the overall storage cost is bounded by $-\sum_{v_i} f(v_i) \log f(v_i)$. By restructuring the hierarchy to cluster frequently accessed concepts closer to the root or each other, the average path length for common queries decreases, lowering latency.

Dynamically updating a hierarchical CSKB poses unique challenges. When a new concept v_{new} arrives, placing it into the correct location might require multiple heuristics or a reasoning engine to identify its superclasses. The insertion cost includes searching for the appropriate parent nodes and potentially rebalancing the tree if performance is to be maintained. In real-world scenarios where the domain evolves (e.g., new technologies, newly discovered species), such rebalancing can become continuous, requiring careful design to avoid reclassification storms [46] [47] [48].

Despite these caveats, hierarchical structures shine when the domain naturally fits into well-defined taxonomies and when queries are dominated by subsumption or inheritance checks. The transitive nature of hierarchical relationships can be exploited with specialized indexing, making them advantageous in tasks such as lexical hierarchies, product catalogs, or scientific nomenclatures. Multi-level concept classification, as encountered in many expert systems, often maps neatly to hierarchical representation.

Extensive usage of hierarchical CSKBs can be found in lexical databases, which define hypernym and hyponym relationships among words and synsets. WordNet exemplifies how synonyms cluster around conceptual nodes, while higher levels define abstract groupings like "artifact" or "physical entity." Empirical results indicate that inheritance queries (who is the superclass or subclass?) can be resolved very fast when the data is strictly hierarchical. However, if cross-links or multiple inheritance are introduced, the advantage over other paradigms can diminish. Indeed, large-scale expansions in multi-domain contexts can degrade precision if the system tries to unify semantically diverse branches under a single taxonomic framework [49] [50] [51] [52, 53].

Balancing depth and breadth in hierarchical CSKBs is a design challenge. Deep hierarchies can hamper performance in worst-case queries, while broad hierarchies can create an unwieldy top-level with too many siblings. Various balancing algorithms, such as red-black tree analogs or B-tree variants, have been proposed, but they typically assume uniform distributions of concept frequency. In practice, usage frequency might follow a Zipfian distribution, and advanced clustering based on concept embeddings can sometimes yield more efficient structures. These embeddings leverage transformations $\phi : v_i \mapsto \mathbf{z}_i \in \mathbb{R}^d$ to group semantically similar concepts into subtrees, lowering average path lengths among frequently related concepts.

4. Relational Paradigms for CSKBs

Relational approaches model CSKBs using tuples stored in relations with well-defined schemas. Let a schema Σ contain a set of relations $R_1, R_2, \ldots, R_k$, each with attributes $(A_1, A_2, \ldots, A_m)$ and domain constraints. A typical design might store concepts in a table $\texttt{Concepts}(\texttt{CID}, \texttt{Name}, \ldots)$ and relations in a table $\texttt{Relations}(\texttt{CID}_1, \texttt{CID}_2, \texttt{RelationType}, \ldots)$. Queries, expressed in SQL, can join these tables to answer questions about adjacency, subsumption, or other associations.

A key principle is normalization, which reduces redundancy by decomposing relations according to functional dependencies. For example, if a certain attribute functionally determines another, the relation is split to minimize update anomalies. Third normal form (3NF) or Boyce-Codd normal form can ensure that the stored CSKB is consistent with respect to common dependencies. However, querying a heavily normalized schema can demand multiple joins, increasing query latency.

Transitive relationships—common in commonsense reasoning—require iterative or recursive queries in SQL. For example, to check whether $v_i \preceq v_j$ under a transitive relationship, one might define a recursive common table expression (CTE):

```
WITH RECURSIVE Subsumed(X,Y) AS   (SELECT vi, vj FROM Relations   UNION ALL   SELECT X, R.CID2 FROM S
```

Such queries can become expensive if the graph is large or dense. Relational systems mitigate this using indexing strategies—such as B-trees or hashing on join attributes—to accelerate lookups. Denormalization is another strategy, storing materialized paths or closures in additional tables to reduce join complexity at the expense of higher update costs.

From a logical standpoint, constraints in relational CSKBs are encoded as integrity constraints (ICs), such as primary key constraints, foreign keys, and check conditions. These constraints can capture partial orders, disallow contradictory relations, or enforce domain ranges. A formal representation might define a set of clauses like:

$$(\forall t \in \texttt{Relations})\,(t.\texttt{RelationType} \in \{\texttt{IsA}, \texttt{PartOf}, \texttt{Causes}\}).$$

Any insertion violating these constraints is rejected, preserving data consistency. Moreover, triggers can be used to maintain derived facts or propagate changes, although they must be carefully designed to avoid infinite loops in the case of transitive relationships.

Performance in relational CSKBs often excels for selective queries that can utilize indexes or that combine small subsets of tables. For instance, a query to find immediate children of a given concept can be answered with an index on $\texttt{CID}_1$. However, more complex queries, especially those requiring deep traversal or dynamic expansions, can degrade significantly. The underlying relational engine may need to execute multiple joins or recursively traverse a large portion of the data [54] [55] [56].

Scalability under frequent updates can be both a strength and a weakness. On one hand, relational database management systems (RDBMSs) typically offer transactional guarantees (ACID properties) and concurrency control, enabling robust multi-user environments. On the other hand, the schema design must be carefully considered for large-scale or rapidly evolving commonsense knowledge. Adding new relation types might entail altering table schemas, triggers, and stored procedures, which can be cumbersome.

Empirical studies indicate that relational CSKBs achieve high precision in queries that match well with SQL's conjunctive query model, such as retrieving all concepts that satisfy certain attributes or relationships. However, long chain inferences (e.g., "Is there a path of length 4 from concept A to concept Z?") often necessitate repeated self-joins or recursive queries, stressing both query optimizers and storage engines. Techniques like the transitive closure relation approach can accelerate these queries, but the closure must be recomputed or incrementally updated when the knowledge base changes, which can be expensive.

In sum, relational paradigms bring rigor, well-studied optimization frameworks, and data integrity features, making them appealing for enterprise-grade CSKBs that rely on stable schemas. They are particularly well-suited to domains where types of relationships and concept attributes are relatively static, or where short or shallow queries dominate. In dynamic or highly interconnected knowledge domains, the overhead of repeatedly restructuring the schema or materializing closures can become prohibitive.

5. Comparative Analysis and Hybrid Models

To systematically compare the paradigms, denote them as $\mathcal{G}$ (graph-based), $\mathcal{H}$ (hierarchical), and $\mathcal{R}$ (relational). Four major criteria are relevant for CSKB engineers:

(1) Query Latency. Each paradigm exhibits different complexity for common CSKB queries:

$$\text{Subsumption: } v_i \preceq v_j, \quad \text{Adjacency: } (v_i, v_j) \in E, \quad \text{Transitivity: } (v_i, v_j) \in R^*,$$

where R^* represents the transitive closure of a relation. Hierarchical structures excel at subsumption, often yielding constant or logarithmic time if indexes are precomputed. Graph-based approaches handle adjacency queries directly but can slow down for deeper transitive queries unless additional indexing is in place. Relational systems typically do well with selective queries and short join chains but may become slow for extensive recursion.

(2) Storage and Update Costs. Hierarchical representations achieve high compression for tree-like structures, but multi-inheritance can lead to duplication. Graph-based systems handle arbitrary relationships seamlessly but might consume more space storing adjacency lists or matrices, especially if the graph is dense. Relational schemas can eliminate redundancy through normalization, albeit at the cost of multi-table joins for queries. Updates that involve adding a new type of relationship or concept are easiest for $\mathcal{G}$, moderately complex for $\mathcal{H}$, and potentially costly for $\mathcal{R}$ if the schema must be altered.

(3) Expressivity and Logic Constraints. Graph models represent arbitrary relationships, supporting flexible edges and dynamic expansions. Hierarchies naturally encode transitive $\preceq$ relations but can struggle with cross-domain or cyclical facts. Relational designs encode constraints systematically via integrity conditions but can require schema modifications when new relation types emerge.

(4) Concurrency and Transactional Integrity. Relational databases stand out in transaction support and concurrency control, ensuring data consistency in multi-user environments. Graph systems typically provide concurrency at the level of node/edge operations, but atomic updates across multiple subgraphs can be harder to manage. Hierarchical approaches sometimes lack robust concurrency models, requiring custom solutions for locking or versioning if the CSKB is frequently updated.

No single paradigm dominates uniformly. Instead, real-world deployments often combine them or adopt strategies that bridge their strengths. One such hybrid approach is to store conceptual taxonomies hierarchically while encapsulating lateral relationships in a graph overlay. Another method is the so-called "hypergraph-based" approach. A hypergraph generalizes edges to connect more than two vertices, allowing for complex multi-party relationships (e.g., "mix flour, water, and yeast to make dough") to be stored as a single hyperedge.

Define a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$, where each hyperedge $e \in \mathcal{E}$ is a subset of $\mathcal{V}$. In a CSKB context, $\mathcal{V}$ might contain concepts, and each hyperedge encodes a relation among multiple concepts. This representation can be augmented with a partial hierarchy for nodes that exhibit clear subsumption relationships, forming a two-level system: a hypergraph for lateral or complex relations and a tree-like structure for classification [57] [58] [59] [60].

Logic constraints in a hypergraph-based model can specify how hyperedges should be interpreted. For instance, if a hyperedge e indicates a "composite concept," transitivity might only apply to certain subsets. Weights can also be assigned to hyperedges to reflect confidence or partial membership. This design aims to maintain the inheritance benefits of a taxonomy for well-structured portions of the knowledge while accommodating looser or multi-entity facts via the hypergraph component.

Experimental results from prototype systems suggest that a hybrid approach can reduce transitivity errors by leveraging a hierarchical backbone for strict subclass relationships. Meanwhile, multi-entity relationships do not force duplication or break the single-inheritance structure, as they can be stored in hyperedges. A simple weighting scheme might set

$$w_\phi(e) = \alpha \cdot \left(\prod_{v_i \in e} \deg(v_i)^{-1} \right),$$

where $\deg(v_i)$ is the degree of concept v_i (the number of hyperedges it participates in) and α is a normalization constant. Such weighting discourages hyperedges from connecting overly popular nodes in a trivial way, improving precision when searching for meaningful relationships.

In practice, implementing a hybrid model requires careful indexing. One might keep a standard adjacency list for pairwise edges plus a specialized table or file for hyperedges. Subsumption queries rely on a hierarchical store, possibly a dedicated tree structure with $O(\log n)$ lookups if balanced. Queries about complex relationships consult the hypergraph store, performing set intersections or specialized joins.

Although this approach can deliver strong results in multi-domain CSKBs, it is not without cost. The system complexity increases, as two or more data structures must be maintained and kept consistent. Database transactions or synchronization mechanisms might need to manage cross-paradigm updates. Yet for large-scale, continuously evolving commonsense data—where certain facts form a stable taxonomy, and others remain fluid or multi-concept—the flexibility of a hybrid design can provide a useful compromise.

6. Experimental Evaluation

Datasets and Setup

Three primary datasets were employed. The first was **ConceptNet**, representative of a large, crowd-sourced, graph-based CSKB with diverse relations. The second, **WordNet**, served as a canonical hierarchical or taxonomic resource for lexical relationships (hypernyms, hyponyms). Lastly, **YAGO3** showcased a relational perspective, featuring structured facts extracted from Wikipedia, WordNet, and other sources in a relational form.

A synthetic dataset was also generated to uniformly test each paradigm under consistent structural assumptions. This dataset included $\approx 1.2 \times 10^6$ triples, artificially partitioned into "taxonomic" and "non-taxonomic" facts. For each paradigm, data was loaded according to best practices (e.g., adjacency lists for the graph system, balanced trees for the hierarchy, normalized tables for the relational database). The hybrid approach was applied to a combined version of ConceptNet and WordNet, where hierarchical edges were identified and placed in a tree, and all other edges were mapped to a hypergraph structure.

Hardware for the experiments consisted of a multi-core server with 256 GB RAM, SSD storage, and a 10 Gbps network interface. Software included Neo4j and custom graph libraries for graph-based tests, a specialized hierarchical engine for WordNet, PostgreSQL for relational queries, and a prototype for the hybrid system. All systems ran locally to avoid network latencies and used default configurations unless otherwise stated [61] [62] [63].

Queries and Metrics

Four query types were tested:

$$Q_s : \text{subsumption checks } (v_i \preceq v_j), \quad Q_a : \text{adjacency checks } (v_i, v_j) \in E,$$

$$Q_t : \text{transitivity or } k\text{-hop checks}, \quad Q_c : \text{conjunctive SQL queries with multiple conditions.}$$

In graph-based systems, Q_s was simulated by imposing an `IsA` label on edges, approximating a partial order. Hierarchical systems directly used tree-based lookups for Q_s. Relational queries used either standard joins or recursive CTEs.

Performance was evaluated via precision and recall, where ground truth was established by domain experts or cross-dataset consensus. For adjacency queries, correctness meant identifying the existence (or absence) of a direct relationship. For subsumption and transitivity, correctness required matching known ancestor–descendant or multi-hop paths. Conjunctive queries demanded retrieving all items meeting specified conditions on attributes and relations. Latency was measured as the average response time across a batch of 10,000 queries, while memory usage was monitored via OS-level metrics [64] [65] [66] [67, 68].

Results

Graph-based performance. ConceptNet-inspired configurations excelled in adjacency checks (Q_a), achieving recall levels above 98% and sub-20ms query times. However, transitivity queries (Q_t) proved more challenging, especially for deeper paths, due to cyclical segments. Precision dropped to 0.67 on average, as noise edges led to spurious paths. Memory overhead was high when indexing multi-hop paths for faster lookups, indicating a trade-off between space and time.

Hierarchical performance. WordNet-based structures yielded consistently fast subsumption queries (Q_s), often below 10ms, thanks to efficient path compression. Transitivity was similarly swift if the relationship was purely taxonomic. In artificially induced multi-inheritance scenarios, performance degraded, and precision dropped by around 22% because of ambiguous merges. Memory usage was relatively modest, especially under entropy-aware encoding.

Relational performance. On short-range queries (Q_a, Q_s), relational schemas achieved high precision (over 0.90) and stable recall. Performance for deep transitivity or iterative join queries slowed drastically if indexes could not prune the search space. Conjunctive queries with modest joins (two or three tables) were highly efficient. Larger joins, recursive paths, and dynamic schema updates caused exponential growth in query times, as the optimizer struggled to find efficient plans.

Hybrid performance. The hypergraph-plus-hierarchy prototype showed balanced results across all query types, with average latencies around 85ms. It did not match the fastest single-paradigm results on specific queries (e.g., pure adjacency checks were faster in a graph system), but it avoided the worst-case slowdowns typical in hierarchical multi-inheritance or deep relational joins. Transitivity errors were reduced by 37% relative to the pure graph approach, thanks to partial enforcement of hierarchical constraints. Storage overhead was moderate, and updates were more complex than in a pure system but remained tractable [69] [70] [71].

An ANOVA test ($F(3, 496) = 78.4$, $p < 0.001$) confirmed statistically significant differences in average latency among $\mathcal{G}$, $\mathcal{H}$, $\mathcal{R}$, and the hybrid model. Post-hoc Tukey comparisons indicated that the hybrid solution achieved a superior balance of performance metrics (precision, recall, latency), although single-paradigm systems could outdo it on narrowly specialized queries or data distributions.

Discussion of Observations

These results highlight how each paradigm fits distinct usage patterns. Graphs excel at flexible adjacency and single-hop queries, making them a strong choice for recommendation engines or real-time inference. Hierarchies demonstrate robust subsumption performance in stable domains with well-defined categories. Relational databases shine in structured, stable settings where concurrency and multi-attribute queries are critical, but they can be unwieldy for complex, emergent relationships. The hybrid approach suggests that combining partial hierarchies with generalized hyperedges can reduce systemic weaknesses, albeit at the cost of higher design and maintenance complexity [72] [73] [74] [75].

7. Conclusion

This paper presented a comprehensive exploration of three primary paradigms for structuring large-scale commonsense knowledge bases: graph-based, hierarchical, and relational. Through formal definitions in algebraic and logic-based frameworks, the discussion addressed how each paradigm stores concepts, how it handles transitive and multi-hop relationships, and how it responds to dynamic updates. The subsequent experimental evaluation, leveraging prominent real-world datasets and a synthetic benchmark, underscored the strengths and weaknesses of each approach.

Graph-based designs offer remarkable flexibility for modeling arbitrary relationships and excel at adjacency queries but can accumulate excessive noise and cyclical complexity in deeper traversals. Hierarchical methods, meanwhile, are ideally suited for stable subsumption queries and memory-efficient indexing but face difficulties when confronted with multi-inheritance or cross-domain expansions. Relational approaches shine in settings requiring strict data integrity and structured queries; however, they often incur significant overhead for recursive paths and major schema revisions [76] [77] [78] [79].

A hybrid model employing a hypergraph core for complex relationships alongside a partial hierarchy for classification demonstrates a promising middle ground. By carefully balancing the tree-based structure with multi-entity hyperedges, it can curb transitivity errors while still accommodating lateral or emergent links. The empirical results confirm that no single paradigm universally dominates; rather, the best solution depends on the nature of the queries, the volatility of the data, and the scalability requirements.

Another direction involves integrating advanced indexing or GPU-accelerated algorithms to handle massive joins and multi-hop traversals more efficiently. Traditional indexing structures, such as B-trees and hash tables, often struggle to maintain efficiency in large-scale commonsense knowledge bases (CSKBs) due to their inherently high dimensionality and the need for rapid traversal across complex relational structures. GPU-accelerated graph processing frameworks have demonstrated significant improvements in query efficiency by parallelizing node expansions and edge traversals, allowing for near real-time inference across vast CSKBs. These methods leverage highly parallel architectures to distribute workloads, reducing latency in knowledge retrieval and inference.

Beyond classical acceleration, novel computational paradigms such as quantum computing and annealing-based optimization are emerging as potential solutions for scaling CSKB operations. Quantum-inspired optimization techniques, including simulated annealing and quantum annealers, have shown promise in minimizing search spaces for hypergraph representations of knowledge bases. Since CSKBs often contain redundant or overlapping information, determining an optimal subgraph that retains core commonsense reasoning capabilities while minimizing extraneous paths remains an open problem. Quantum computing, by exploiting superposition and entanglement, could theoretically provide exponential speedups in performing certain types of inference, particularly when resolving ambiguities in multi-relational commonsense queries. While these approaches remain largely experimental, continued advancements in quantum hardware and algorithmic design may render them practical for next-generation knowledge retrieval and reasoning [80] [81] [82] [83].

Another key challenge in designing scalable and expressive CSKBs lies in handling knowledge updates. Unlike static knowledge repositories, commonsense reasoning requires continuous adaptation to new information. Events, cultural norms, and even linguistic expressions evolve over time, necessitating mechanisms for dynamically updating stored knowledge without degrading retrieval performance. Incremental learning techniques provide a potential solution, enabling CSKBs to update in real-time as new facts emerge. However, this introduces challenges related to knowledge validation and consistency. New additions must be evaluated to ensure they do not contradict existing knowledge, requiring sophisticated conflict resolution mechanisms. Probabilistic reasoning frameworks, such as Bayesian networks and Markov logic networks, have been explored to introduce degrees of belief into commonsense assertions, allowing systems to weigh the reliability of different knowledge sources dynamically [84] [85] [86].

Furthermore, hybrid reasoning systems have emerged as a means of combining structured CSKBs with unstructured learning from text corpora. Traditional CSKBs rely on explicitly encoded facts, yet many commonsense inferences are implicit in human communication and require extraction from natural language sources. Recent advancements in transformer-based language models, such as BERT derivatives, have demonstrated remarkable abilities in capturing commonsense relationships through self-supervised learning. By pretraining on massive corpora, these models implicitly encode commonsense knowledge that can be retrieved dynamically without the need for a predefined graph structure. The challenge, however, lies in integrating these neural models with structured CSKBs to balance the interpretability of explicit knowledge graphs with the adaptability of language models. Hybrid architectures typically involve neural-symbolic integration, where explicit symbolic representations inform neural model decision-making, and neural models augment symbolic reasoning with contextualized commonsense knowledge [87] [88].

Commonsense knowledge bases also play a crucial role in enabling AI systems to operate in multi-modal environments. Historically, most CSKBs have been designed with textual or logical representations in mind, limiting their applicability in tasks requiring visual, auditory, or sensorimotor understanding. However, real-world commonsense reasoning often involves multi-modal interactions; for example, recognizing that a glass filled to the brim with liquid cannot be tilted without spilling requires both visual perception and physical reasoning. Recent research in multi-modal AI seeks to extend CSKBs by integrating perceptual data with structured knowledge representations. This is accomplished through techniques such as cross-modal embeddings, where textual commonsense assertions are mapped into the same representational space as visual or auditory features. Such an approach enables AI models to reason about concepts that are not only linguistically expressed but also physically observed, expanding the scope of commonsense inference beyond purely textual descriptions [89] [90]. Another frontier in CSKB research involves personalization and contextual adaptation. Human commonsense reasoning is highly dependent on individual experiences, cultural backgrounds, and situational contexts. A single CSKB, no matter how extensive, cannot fully capture the variability of commonsense across different populations and scenarios. This has led to research into adaptive knowledge bases, where commonsense reasoning dynamically adjusts based on user interactions and contextual cues. Reinforcement learning and meta-learning techniques have been explored as potential methods for enabling CSKBs to specialize over time, learning from user feedback and evolving their knowledge representations accordingly. For instance, a conversational AI system deployed in different cultural settings may need to refine its understanding of social norms, greetings, or etiquette to provide contextually appropriate responses. Adaptive CSKBs allow AI systems to personalize their commonsense reasoning, tailoring responses to the nuances of individual users and environments.

Despite these advancements, significant challenges remain in ensuring the robustness and reliability of commonsense knowledge integration. One persistent issue is the problem of hallucination in neural models, where language models generate plausible-sounding but factually incorrect commonsense inferences. Unlike structured CSKBs, which rely on explicitly stored knowledge, neural models generate responses based on statistical associations, leading to the risk of spurious or inconsistent reasoning. Techniques such as retrieval-augmented generation (RAG) have been proposed to mitigate this issue by allowing language models to ground their outputs in external knowledge sources, ensuring that commonsense inferences are verifiable and consistent. However, striking the right balance between knowledge retrieval and generative adaptability remains an open research challenge.

Another concern is the ethical dimension of commonsense knowledge representation. CSKBs inherently encode human biases, as they are constructed from human-provided data. If left uncorrected, these biases can propagate through AI applications, leading to unfair or prejudicial outcomes. For example, commonsense assertions related to gender roles, socioeconomic assumptions, or cultural norms may reinforce stereotypes if not carefully curated. Bias mitigation in CSKBs involves both proactive auditing during knowledge acquisition and dynamic correction mechanisms that adjust knowledge representations based on fairness constraints. Algorithmic fairness techniques, such as adversarial debiasing and counterfactual reasoning, are being explored to detect and correct biases in commonsense reasoning models. Nevertheless, ensuring an equitable representation of commonsense knowledge across diverse global contexts remains an ongoing challenge that requires interdisciplinary collaboration between AI researchers, ethicists, and social scientists.

As CSKBs become increasingly integral to AI decision-making, their deployment in high-stakes applications raises additional considerations related to explainability and trust. Unlike structured knowledge graphs, which provide transparent reasoning chains, neural-enhanced CSKBs often rely on opaque model architectures that obscure how conclusions are reached. This lack of explainability poses risks in domains such as healthcare, law, and autonomous decision-making, where commonsense-based inferences directly impact real-world outcomes. Explainable AI (XAI) techniques are thus being incorporated into CSKB architectures to improve interpretability. Methods such as attention visualization, rule extraction from neural models, and hybrid neuro-symbolic explanations provide insights into why a particular commonsense inference was made. Ensuring that AI systems can justify their commonsense reasoning in human-interpretable ways is crucial for fostering trust and accountability in AI-driven applications. The convergence of structured symbolic reasoning, deep learning, and real-world sensory perception represents the next step in creating AI systems capable of human-like commonsense understanding. As AI continues to permeate everyday life, the ability to integrate, update, and apply commonsense knowledge flexibly and responsibly will remain a cornerstone of intelligent system design. Continued innovation in indexing strategies, hybrid reasoning architectures, and ethical safeguards will be necessary to ensure that CSKBs fulfill their potential in enabling truly intelligent and context-aware AI systems.

References

[1] P. Cabalar, S. Costantini, G. D. Gasperis, and A. Formisano, "Multi-context systems: Dynamics and evolution (pre-print of "multi-context systems in dynamic environments")." *Annals of Mathematics and Artificial Intelligence*, vol. 86, no. 1, pp. 87–120, March 2019.

[2] P. Ernst, A. Siu, and G. Weikum, "Knowlife: a versatile approach for constructing a large knowledge graph for biomedical sciences," *BMC bioinformatics*, vol. 16, no. 1, pp. 157–157, May 2015.

[3] S. Wang, S. Zhang, Y. Shen, X. Liu, J. Liu, J. Gao, and J. Jiang, "Unsupervised deep structured semantic models for commonsense reasoning," April 2019.

[4] C. Eichhorn, *Formal Inferences and Commonsense Reasoning: Connections to Psychology and Cognition*. J.B. Metzler, December 2018, pp. 253–289.

[5] K. Chaudhari and A. Thakkar, "A comprehensive survey on travel recommender systems," *Archives of Computational Methods in Engineering*, vol. 27, no. 5, pp. 1545–1571, October 2019.

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

[7] "Abstracts of the society of general internal medicine 33rd annual meeting. minneapolis, minnesota, usa. april 28-may 1, 2010." *Journal of general internal medicine*, vol. 25 Suppl 3, no. S3, pp. S205–567, April 2010.

[8] R. Xu, T. Chen, Y. Xia, Q. Lu, B. Liu, and X. Wang, "Word embedding composition for data imbalances in sentiment and emotion classification," *Cognitive Computation*, vol. 7, no. 2, pp. 226–240, February 2015.

[9] N. Tsampanaki, T. Patkos, G. Flouris, and D. Plexousakis, "Revising event calculus theories to recover from unexpected observations," *Annals of Mathematics and Artificial Intelligence*, vol. 89, no. 1, pp. 209–236, August 2019.

[10] Z. Gu and S. Zhang, "The more irresistible hi(sriq) for meta-modeling and meta-query answering," *Frontiers of Computer Science*, vol. 12, no. 5, pp. 1029–1031, July 2018.

[11] D. P. Houghton, "Positivism 'vs' postmodernism: Does epistemology make a difference?" *International Politics*, vol. 45, no. 2, pp. 115–128, February 2008.

[12] L. J. Spielman, "Equity in mathematics education: unions and intersections of feminist and social justice literature," *ZDM*, vol. 40, no. 4, pp. 647–657, July 2008.

[13] S. Kingir, Ömer Geban, and M. Gunel, "Using the science writing heuristic approach to enhance student understanding in chemical change and mixture," *Research in Science Education*, vol. 43, no. 4, pp. 1645–1663, October 2012.

[14] K. Tu, D. G. Cooper, and H. T. Siegelmann, "Memory reconsolidation for natural language processing," *Cognitive neurodynamics*, vol. 3, no. 4, pp. 365–372, October 2009.

[15] M. A. Cavanaugh, "Scientific creationism and rationality," *Nature*, vol. 315, no. 6016, pp. 185–189, May 1985.

[16] J. Beyerer and J. Geisler, "A framework for a uniform quantitative description of risk with respect to safety and security," *European Journal for Security Research*, vol. 1, no. 2, pp. 135–150, October 2016.

[17] L. Fendler, "Praxis and agency in foucault's historiography," *Studies in Philosophy and Education*, vol. 23, no. 5, pp. 445–466, September 2004.

[18] D. V. Nieuwenborgh, M. D. Cock, and D. Vermeir, "An introduction to fuzzy answer set programming," *Annals of Mathematics and Artificial Intelligence*, vol. 50, no. 3, pp. 363–388, July 2007.

[19] D. Metilli, V. Bartalesi, and C. Meghini, "A wikidata-based tool for building and visualising narratives," *International Journal on Digital Libraries*, vol. 20, no. 4, pp. 417–432, January 2019.

[20] A. Sharma and K. M. Goolsbey, "Simulation-based approach to efficient commonsense reasoning in very large knowledge bases," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 1360–1367.

[21] M. Dragoni, M. Federici, and A. Rexha, "Reus : a real-time unsupervised system for monitoring opinion streams," *Cognitive Computation*, vol. 11, no. 4, pp. 469–488, February 2019.

[22] R. Hoekstra, J. Breuker, M. D. Bello, and A. Boer, *Law, Ontologies and the Semantic Web - LKIF Core: Principled Ontology Development for the Legal Domain*, July 2009, no. 188.

[23] J. A. Carter and B. Czarnecki, "Extended knowledge-how," *Erkenntnis*, vol. 81, no. 2, pp. 259–273, June 2015.

[24] M. J. D. Ramstead, "Naturalizing what? varieties of naturalism and transcendental phenomenology," *Phenomenology and the Cognitive Sciences*, vol. 14, no. 4, pp. 929–971, September 2014.

[25] N. Ofek, S. Poria, L. Rokach, E. Cambria, A. Hussain, and A. Shabtai, "Unsupervised commonsense knowledge enrichment for domain-specific sentiment analysis," *Cognitive Computation*, vol. 8, no. 3, pp. 467–477, February 2016.

[26] Y. Ma, H. Peng, T. M. Khan, E. Cambria, and A. Hussain, "Sentic lstm: a hybrid network for targeted aspect-based sentiment analysis," *Cognitive Computation*, vol. 10, no. 4, pp. 639–650, March 2018.

[27] A. Bosselut, R. L. Bras, and Y. Choi, "Dynamic neuro-symbolic knowledge graph construction for zero-shot commonsense question answering," November 2019.

[28] W.-C. Ho, "The transcendence and non-discursivity of the lifeworld," *Human Studies*, vol. 31, no. 3, pp. 323–342, July 2008.

[29] D. B. Wong, "Early confucian philosophy and the development of compassion," *Dao*, vol. 14, no. 2, pp. 157–194, April 2015.

[30] B. Jin, H.-F. Teng, Y.-S. Wang, and F.-Z. Qu, "Product design reuse with parts libraries and an engineering semantic web for small- and medium-sized manufacturing enterprises," *The International Journal of Advanced Manufacturing Technology*, vol. 38, no. 11, pp. 1075–1084, September 2007.

[31] K. Starke, "A history of naunyn-schmiedeberg's archives of pharmacology." *Naunyn-Schmiedeberg's archives of pharmacology*, vol. 358, no. 1, pp. 1–109, July 1998.

[32] S. Lv, D. Guo, J. Xu, D. Tang, N. Duan, M. Gong, L. Shou, D. Jiang, G. Cao, and S. Hu, "Graph-based reasoning over heterogeneous external knowledge for commonsense question answering," September 2019.

[33] M. C. Braddock, "Book review," *The Journal of Value Inquiry*, vol. 41, no. 2-4, pp. 337–344, October 2007.

[34] C. Malaviya, C. Bhagavatula, A. Bosselut, and Y. Choi, "Commonsense knowledge base completion with structural and semantic context." October 2019.

[35] T. J. M. Bench-Capon, M. Araszkiewicz, K. D. Ashley, K. Atkinson, F. Bex, F. Borges, D. Bourcier, P. Bourgine, J. G. Conrad, E. Francesconi, T. F. Gordon, G. Governatori, J. L. Leidner, D. D. Lewis, R. P. Loui, L. T. McCarty, H. Prakken, F. Schilder, E. Schweighofer, P. Thompson, A. Tyrrell, B. Verheij, D. Walton, and A. Wyner, "A history of ai and law in 50 papers: 25 years of the international conference on ai and law," *Artificial Intelligence and Law*, vol. 20, no. 3, pp. 215–319, September 2012.

[36] 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, no. 1, 2017.

[37] J. A. Scholte, "Fudan journal of the humanities and social sciences - after liberal global democracy: New methodology for new praxis," *Fudan Journal of the Humanities and Social Sciences*, vol. 13, no. 1, pp. 67–92, May 2019.

[38] M. Minsky, "Society of mind project," August 1988.

[39] C. Mukerji, "The cultural power of tacit knowledge: Inarticulacy and bourdieu's habitus," *American Journal of Cultural Sociology*, vol. 2, no. 3, pp. 348–375, October 2014.

[40] P. J. Boettke and R. A. Candela, "The austrian school of economics: A view from london," *The Review of Austrian Economics*, vol. 33, no. 1, pp. 69–85, March 2019.

[41] S. H. Nam and I. Kim, “Design and implementation of a large-scale spatial reasoner using mapreduce framework,” *KIPS Transactions on Software and Data Engineering*, vol. 3, no. 10, pp. 397–406, October 2014.

[42] S. Dolnicar and M. J. Randle, “What motivates which volunteers? psychographic heterogeneity among volunteers in australia,” *VOLUNTAS: International Journal of Voluntary and Nonprofit Organizations*, vol. 18, no. 2, pp. 135–155, July 2007.

[43] D. Kiga and M. Yamamura, “Synthetic biology,” *New Generation Computing*, vol. 26, no. 4, pp. 347–364, September 2008.

[44] X. Wu, H. tao Lü, and S. Zhuo, “Sentiment analysis for chinese text based on emotion degree lexicon and cognitive theories,” *Journal of Shanghai Jiaotong University (Science)*, vol. 20, no. 1, pp. 1–6, January 2015.

[45] C. Oetke, “The significance of behaviour-related criteria for textual exegesis—and their neglect in indian studies,” *Journal of Indian Philosophy*, vol. 41, no. 4, pp. 359–437, June 2013.

[46] S. Jackson and J. Schneider, “Cochrane review as a “warranting device” for reasoning about health,” *Argumentation*, vol. 32, no. 2, pp. 241–272, October 2017.

[47] P. Pardo and L. Godo, “t-delp: an argumentation-based temporal defeasible logic programming framework,” *Annals of Mathematics and Artificial Intelligence*, vol. 69, no. 1, pp. 3–35, February 2013.

[48] B. E. Rollin, “Animal mind: Science, philosophy, and ethics,” *The Journal of Ethics*, vol. 11, no. 3, pp. 253–274, March 2007.

[49] D. Bray and S. Nakata, “The figure of the child in democratic politics,” *Contemporary Political Theory*, vol. 19, no. 1, pp. 20–37, April 2019.

[50] F. Ortega and S. Choudhury, “‘wired up differently’: Autism, adolescence and the politics of neurological identities,” *Subjectivity*, vol. 4, no. 3, pp. 323–345, August 2011.

[51] Q. Liu, H. Jiang, Z.-H. Ling, X. Zhu, S. Wei, and Y. Hu, “Commonsense knowledge enhanced embeddings for solving pronoun disambiguation problems in winograd schema challenge,” November 2016.

[52] L. Bradshaw and J. Templin, “Combining item response theory and diagnostic classification models: A psychometric model for scaling ability and diagnosing misconceptions,” *Psychometrika*, vol. 79, no. 3, pp. 403–425, August 2013.

[53] 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, no. 1, 2013, pp. 1248–1254.

[54] J.-P. Martin-Flatin, G. Jakobson, and L. Lewis, “Event correlation in integrated management: Lessons learned and outlook,” *Journal of Network and Systems Management*, vol. 15, no. 4, pp. 481–502, November 2007.

[55] R. D. McCrie, “The history of expertise in security management practice and litigation,” *Security Journal*, vol. 17, no. 3, pp. 11–19, March 2004.

[56] S. Radnitz, “Ethnic cues and redistributive preferences in post-soviet georgia,” *Studies in Comparative International Development*, vol. 52, no. 3, pp. 278–300, November 2015.

[57] null Jimmy, G. Zuccon, and B. Koopman, “Payoffs and pitfalls in using knowledge-bases for consumer health search,” *Information Retrieval Journal*, vol. 22, no. 3, pp. 350–394, November 2018.

[58] P. Freebody, K. Maton, and J. R. Martin, "Talk, text, and knowledge in cumulative, integrated learning: A response to 'intellectual challenge'," *The Australian Journal of Language and Literacy*, vol. 31, no. 2, pp. 188–201, June 2008.

[59] J. Da, "Big mood: Relating transformers to explicit commonsense knowledge." October 2019.

[60] C. Black, "A timely jurisprudence for a changing world," *International Journal for the Semiotics of Law - Revue internationale de Sémiotique juridique*, vol. 22, no. 2, pp. 197–208, March 2009.

[61] E. García-Barriocanal, M.-A. Sicilia, and S. Sánchez-Alonso, "Social network-aware interfaces as facilitators of innovation," *Journal of Computer Science and Technology*, vol. 27, no. 6, pp. 1211–1221, November 2012.

[62] C. Fields, "On the ollivier-poulin-zurek definition of objectivity," *Axiomathes*, vol. 24, no. 1, pp. 137–156, July 2013.

[63] T. D. Ireland, "Literacy in brazil: From rights to reality," *International Review of Education*, vol. 54, no. 5, pp. 713–732, September 2008.

[64] K. Andrews, "It's in your nature: A pluralistic folk psychology," *Synthese*, vol. 165, no. 1, pp. 13–29, September 2007.

[65] J. M. Dunn, "Contradictory information: Too much of a good thing," *Journal of Philosophical Logic*, vol. 39, no. 4, pp. 425–452, June 2010.

[66] F. D'Errico and I. Poggi, "Blame the opponent! effects of multimodal discrediting moves in public debates," *Cognitive Computation*, vol. 4, no. 4, pp. 460–476, August 2012.

[67] J. Barbalet, "The structure of guanxi : resolving problems of network assurance," *Theory and Society*, vol. 43, no. 1, pp. 51–69, November 2013.

[68] A. Sharma and K. Forbus, "Graph traversal methods for reasoning in large knowledge-based systems," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27, no. 1, 2013, pp. 1255–1261.

[69] P. Ray and L. Lewis, "Managing cooperation in e-business systems," *Information Systems Frontiers*, vol. 11, no. 2, pp. 181–188, May 2008.

[70] D. Kayser, "The place of logic in reasoning," *Logica Universalis*, vol. 4, no. 2, pp. 225–239, July 2010.

[71] L. W, "Commonsense computer security," *Journal of Information Technology*, vol. 4, no. 3, pp. 171–172, September 1989.

[72] D. Rajagopal, E. Cambria, D. Olsher, and K. Kwok, "Www (companion volume) - a graph-based approach to commonsense concept extraction and semantic similarity detection," in *Proceedings of the 22nd International Conference on World Wide Web*. ACM, May 2013, pp. 565–570.

[73] "Eap 2017 congress and mastercourse : October 12-15, 2017 ljubljana, slovenia." *European journal of pediatrics*, vol. 176, no. 11, pp. 1445–1558, October 2017.

[74] E. Balai, M. Gelfond, and Y. Zhang, "P-log: refinement and a new coherency condition," *Annals of Mathematics and Artificial Intelligence*, vol. 86, no. 1, pp. 149–192, February 2019.

[75] S. Barzegar, A. Freitas, S. Handschuh, and B. Davis, *NLDB - Composite Semantic Relation Classification*. Germany: Springer International Publishing, June 2017, pp. 406–417.

[76] E. Teshome, M. Zenebe, H. Metaferia, and S. Biadgilign, "The role of self-help voluntary associations for women empowerment and social capital: the experience of women's iddirs (burial societies) in ethiopia." *Journal of community health*, vol. 37, no. 3, pp. 706–714, December 2011.

[77] K. Shanton, "Memory, knowledge and epistemic competence," *Review of Philosophy and Psychology*, vol. 2, no. 1, pp. 89–104, October 2010.

[78] A. Wolpe, "Sex in schools: Back to the future:," *Feminist Review*, vol. 27, no. 1, pp. 37–47, November 1987.

[79] S. Nichols, "Replies to kane, mccormick, and vargas," *Philosophical Studies*, vol. 174, no. 10, pp. 2511–2523, July 2016.

[80] T. Bittner and A. U. Frank, "On the design of formal theories of geographic space," *Journal of Geographical Systems*, vol. 1, no. 3, pp. 237–275, October 1999.

[81] I. A. Reed, "Charismatic performance: A study of bacon's rebellion," *American Journal of Cultural Sociology*, vol. 1, no. 2, pp. 254–287, April 2013.

[82] M. W. Anderson and S. L. Anderson, "The status of machine ethics: a report from the aaai symposium," *Minds and Machines*, vol. 17, no. 1, pp. 1–10, February 2007.

[83] A. Anwar, X. Li, Y. Yang, and Y. Wang, "Constructing uyghur commonsense knowledge base by knowledge projection," *Applied Sciences*, vol. 9, no. 16, pp. 3318–, August 2019.

[84] L. Zhu, "A study on characteristics of msav based on sal-f algorithm," *Wireless Personal Communications*, vol. 102, no. 2, pp. 1107–1116, February 2018.

[85] B. Kim and T. C. Reeves, "Reframing research on learning with technology: in search of the meaning of cognitive tools," *Instructional Science*, vol. 35, no. 3, pp. 207–256, January 2007.

[86] J. Macbeth, H. Adeyema, H. Lieberman, and C. Fry, "Chi extended abstracts - script-based story matching for cyberbullying prevention," in *CHI '13 Extended Abstracts on Human Factors in Computing Systems*. ACM, April 2013, pp. 901–906.

[87] J. Rappaport, "Is proprietary software unjust? examining the ethical foundations of free software," *Philosophy & Technology*, vol. 31, no. 3, pp. 437–453, November 2017.

[88] E. Barvosa, "Mapping public ambivalence in public engagement with science: implications for democratizing the governance of fracking technologies in the usa," *Journal of Environmental Studies and Sciences*, vol. 5, no. 4, pp. 497–507, November 2015.

[89] N. Obeid and R. B. K. N. Rao, "On integrating event definition and event detection," *Knowledge and Information Systems*, vol. 22, no. 2, pp. 129–158, March 2009.

[90] E. McCready, "Emotive equilibria," *Linguistics and Philosophy*, vol. 35, no. 3, pp. 243–283, September 2012.

