

Commonsense Knowledge and Its Representation in Artificial Intelligence: A Conceptual Review of Symbolic, Graph-Based, Neural, and Neuro-Symbolic Approaches

Ndiana Okon Asuquo¹, Obonguko, Ubon Asuquo²

¹Department of Computing, Top Faith University, Mkpatak, Nigeria

²Department of Computing, Akwa Ibom State Polytechnic, Ikot Osurua, Nigeria

Received: 05.09.2026 | Accepted: 07.10.2026 | Published: 18.10.2026

*Corresponding author: Ndiana Okon Asuquo

DOI: [10.5281/zenodo.23446028](https://doi.org/10.5281/zenodo.23446028)

Abstract

Original Research

Commonsense knowledge comprises the everyday facts, expectations, causal assumptions, social conventions, Physical regularities, and practical inference patterns that people ordinarily use without explicitly stating them. Although artificial intelligence (AI) has achieved substantial progress in natural-language processing, vision, planning, decision support, and generative language, reliable commonsense remains difficult to acquire, represent, retrieve, update, and reason over. This article presents a conceptual review of commonsense knowledge and its representation in AI, synthesizing the major representational traditions from symbolic logic, frames, scripts, and semantic networks to ontologies, knowledge graphs, probabilistic models, neural language models, and neuro-symbolic architectures. Particular attention is given to Cyc, ConceptNet, ATOMIC, and ATOMIC 2020 as influential resources that illustrate complementary approaches to explicit and inferential commonsense. The review shows that commonsense is characterized by implicitness, context dependence, defeasibility, uncertainty, temporal and spatial structure, social situatedness, and cultural variation. These characteristics make a single universal representation difficult to sustain. The article therefore proposes a layered hybrid architecture comprising input interpretation, an explicit commonsense knowledge graph, neural retrieval, rule-based and probabilistic reasoning, and an explanation/provenance layer. The framework is positioned as a research design rather than an experimentally validated system. The article also identifies knowledge acquisition, coverage, inconsistency, temporal change, cultural variation, grounding, scalability, evaluation, provenance, and hallucination control as continuing challenges. Finally, it highlights opportunities for educational AI and context-sensitive systems in Nigerian and African settings, where local linguistic and institutional knowledge can complement globally developed commonsense resources.

Keywords: commonsense knowledge; knowledge representation; artificial intelligence; knowledge graphs; symbolic reasoning; neural reasoning; neuro-symbolic AI; ConceptNet; Cyc; ATOMIC.

Copyright © 2026 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0).

1.1 Background and problem context

Artificial intelligence is concerned with constructing computational systems that perceive environments, represent knowledge,

learn from experience, reason about alternatives, and act toward goals. Knowledge representation and reasoning therefore remain central to AI because intelligent behaviour depends not only

on pattern recognition but also on the ability to connect observations and language with structured accounts of entities, events, relationships, constraints, and possible consequences (Davis, 1990; Russell & Norvig, 2021).

Commonsense knowledge is a particularly demanding form of world knowledge because much of it is implicit. A human reader encountering a statement such as “the lecturer dropped the glass because it was slippery” can ordinarily supply background assumptions about physical objects, gravity, causation, agency, temporal order, and likely consequences without those assumptions being explicitly stated. The difficulty for AI is not simply the size of the required knowledge base. Commonsense includes defeasible expectations, contextual interpretations, qualitative spatial and temporal relations, causal and social knowledge, and reasoning about agents whose goals and beliefs may be only partially observable (Davis & Marcus, 2015).

Recent AI systems can learn broad statistical regularities from very large corpora and can produce fluent language, but statistical plausibility does not by itself guarantee grounded or logically supported commonsense inference. Knowledge graphs and external knowledge resources can supply explicit structure, while neural models offer flexible semantic matching and language interaction. The research problem is therefore increasingly one of integration: how can different forms of representation be selected and combined so that an AI system can use relevant knowledge, revise conclusions when circumstances change, represent uncertainty, and expose evidence for its inferences?

1.2 Aim and objectives

This article aims to synthesize the principal approaches to representing commonsense knowledge in AI and to develop a conceptual framework for combining explicit and learned representations. Specifically, it seeks to:

- clarify the meaning and major dimensions of commonsense knowledge;

- examine the characteristics that make commonsense difficult to represent computationally;
- review symbolic, graphical, probabilistic, neural, and hybrid representation approaches;
- compare major commonsense resources, including Cyc, ConceptNet, ATOMIC, and ATOMIC 2020;
- identify persistent technical and socio-cultural challenges; and
- outline research opportunities for context-aware and educational AI, particularly in Nigerian settings.

1.3 Guiding questions

1. What constitutes commonsense knowledge, and how does it differ from encyclopedic or domain-specific knowledge?
2. How do different representational paradigms capture different dimensions of commonsense?
3. How can explicit knowledge and learned statistical representations be combined?
4. What technical and socio-cultural factors limit reliable commonsense reasoning?
5. What research opportunities arise for commonsense-enabled educational and context-aware AI?

2. Review Approach

This article is a conceptual and narrative review derived primarily from the doctoral seminar manuscript on commonsense knowledge and its representation in AI. The source material was reorganized around the problem definition, characteristics of commonsense, historical representation paradigms, major resources, contemporary reasoning, limitations, and research opportunities. Selected bibliographic records in the seminar reference list were checked against publisher or scholarly-index records where available, including the AAAI records for ConceptNet 5.5, ATOMIC, and ATOMIC 2020 and the IEEE record for the

knowledge-graph survey. The review does not claim to be a systematic review or meta-analysis, and no new experimental dataset, benchmark, or statistical result is introduced. Its principal contribution is therefore conceptual synthesis and the formulation of a layered research architecture.

3. Commonsense Knowledge: Meaning, Scope, and Characteristics

3.1 Meaning and boundaries

Commonsense knowledge can be understood as generally shared practical knowledge used to

interpret ordinary situations and guide appropriate action. It covers objects and their properties, typical actions, physical regularities, causal relationships, intentions, goals, time, space, social interaction, and ordinary expectations (Davis, 1990). The term “shared,” however, should not be treated as synonymous with “universal.” Social conventions, educational practices, linguistic conventions, and interpretations of events can vary across communities. A computational representation should consequently preserve context and provenance rather than encode every socially situated proposition as an unconditional universal truth.

Knowledge type	Typical content	Illustrative example	Representation implication
Commonsense	Everyday expectations and practical inference	A fragile glass can break when dropped.	Defeasible and contextual inference
Encyclopedic	Explicit facts about entities and events	Paris is the capital of France.	Fact retrieval may be sufficient
Domain-specific	Specialized professional knowledge	A compiler transforms source code.	Domain ontology and terminology
Procedural	Actions, states, and task sequences	Steps for submitting an assignment.	Temporal and action structure
Social/affective	Intentions, norms, and reactions	An apology may reduce tension.	Agent, context, and uncertainty modelling

3.2 Dimensions of commonsense

- Physical commonsense: expectations about objects, movement, collision, containment, material properties, and changes of state.
- Temporal commonsense: duration, ordering, persistence, temporal dependency, and change.
- Causal commonsense: typical causes and effects and the distinction between correlation and plausible causal explanation.
- Social commonsense: goals, beliefs, expectations, obligations, norms, intentions, and reactions of agents.
- Linguistic and pragmatic commonsense: interpretation of words and utterances in context and recognition of implied meaning.
- Spatial commonsense: relations such as inside, above, behind, near, between, and relative position.
- Quantitative commonsense: approximate magnitudes and qualitative comparisons where exact numerical calculation is unnecessary.
- Meta-knowledge: representations of what is known, unknown, believed, assumed, or uncertain.




Figure 1. Dimensions of commonsense knowledge identified in the review.

3.3 Why commonsense is difficult to represent

- **Implicitness:** everyday communication leaves much knowledge unstated, so corpora provide indirect and incomplete evidence.
- **Context dependence:** the same proposition can be useful in one context and misleading in another.
- **Defeasibility:** typical conclusions can be withdrawn when exceptions or new evidence appear.
- **Uncertainty:** many commonsense relations are tendencies rather than universal laws.
- **Cultural and linguistic variation:** social expectations and everyday practices differ across communities and languages.
- **Open-endedness:** the number of relevant entities, situations, events, and possible interactions is effectively unbounded.




Figure 2. Historical development of commonsense representation described in the review.

4. Historical Development of Commonsense Representation

The development of commonsense representation can be viewed as a progression from explicit symbolic modelling toward increasingly data-driven representations, followed by renewed interest in hybrid systems. Early AI emphasized logic, theorem proving,

production systems, semantic networks, and frames. Later work expanded structured knowledge through large hand-built knowledge bases, web-scale resources, and knowledge graphs. Neural language models subsequently demonstrated broad learned coverage, while current research increasingly explores combinations of neural and structured representations.

Period	Dominant orientation	Representative techniques/resources	Contribution
1950s–1960s	Formal reasoning	Logic, theorem proving, search	Explicit computational inference
1970s–1980s	Structured conceptual modelling	Frames, scripts, semantic networks	Stereotyped situations and relations
1980s–1990s	Large-scale hand-built knowledge	Cyc and related ontologies	Broad explicit everyday knowledge
2000s	Web-scale/crowdsourced knowledge	Open knowledge bases and lexical resources	Expanded acquisition
2010s	Commonsense knowledge graphs	ConceptNet, WebChild, ATOMIC	Typed relations and inferential patterns
2020s	Neural and hybrid reasoning	ATOMIC 2020, LLMs, neuro-symbolic systems	Combines learned representations with structure

5. Major Approaches to Representing Commonsense Knowledge

5.1 Logic-based representation

Propositional and first-order logic provide precise languages for objects, properties, relations, and quantified statements. They support deduction, consistency checking, and proof procedures. However, many commonsense propositions are defaults rather than universal laws. A rule such as $\forall x \text{ (Bird}(x) \wedge \text{NormalBird}(x) \rightarrow \text{CanFly}(x))$ illustrates the need for exceptions and therefore motivates non-monotonic, default, or probabilistic extensions.

5.2 Frames and scripts

Frames organize stereotyped situations through slots and values, while scripts represent typical event sequences. A lecture script, for example, can include a lecturer, students, a topic, presentation activities, questions, and an ending. These representations are intuitive and useful for recurring situations but can become brittle when real situations diverge from the expected pattern.

5.3 Semantic networks

Semantic networks model concepts as nodes connected by labelled relations. Their graphical structure supports inheritance and relational

navigation and influenced later ontology and knowledge-graph approaches. Their limitation is that simple edges may not adequately express complex temporal, contextual, probabilistic, or defeasible relations.

5.4 Ontologies and description logics

Ontologies define concepts, classes, properties, and constraints within a domain, while description logics provide formal foundations for reasoning over many ontology structures. They are valuable where semantic precision, interoperability, and consistency are important. A formal ontology, however, does not automatically provide the broad background expectations required for everyday reasoning.

5.5 Knowledge graphs

Knowledge graphs connect entities and concepts through typed relations and can integrate information from multiple sources. Research on knowledge graphs spans representation learning, acquisition, completion, temporal knowledge graphs, and reasoning (Ji et al., 2022). Commonsense knowledge graphs extend this paradigm by emphasizing everyday relations and inferential patterns.

5.6 Probabilistic representation

Bayesian and related probabilistic approaches explicitly model uncertainty and can update beliefs as evidence changes. This is attractive because many everyday conclusions are probabilistic. The principal challenges include acquiring suitable distributions, calibration, and scaling probability models to an open-ended world.

5.7 Neural representation

Neural language models learn distributed representations from very large corpora and provide broad language coverage and flexible interaction. Their limitation for commonsense is that statistical plausibility is not equivalent to grounded truth or valid inference. External structured knowledge can therefore serve as a source of retrieval, constraints, or evidence.

5.8 Neuro-symbolic representation

Neuro-symbolic AI combines learned pattern recognition with explicit representations and reasoning procedures. A neural component may interpret language and retrieve relevant concepts, while symbolic or graph-based components can apply constraints, check relations, and expose an inference path. This combination is particularly relevant because commonsense simultaneously demands broad learning, explicit structure, contextual interpretation, and reasoning.

6. Major Commonsense Knowledge Resources

Several resources illustrate different assumptions about what should be represented and how commonsense should be organized. They should not be treated as interchangeable: Cyc emphasizes explicit axioms and logical inference; ConceptNet emphasizes general relational knowledge; ATOMIC emphasizes event-centred inferential knowledge; and ATOMIC 2020 extends the commonsense knowledge-graph direction toward integration with neural knowledge models.


Resource	Primary orientation	Representation	Main contribution	Key limitation
Cyc	Broad everyday knowledge	Axioms, predicates, formal knowledge	Deep explicit knowledge and logical inference	Knowledge engineering and maintenance cost

ConceptNet	General commonsense relations	Labelled multilingual graph	Large relational coverage and accessible structure	Assertion quality and context vary
ATOMIC	Event and social inference	Typed if-then relations	Causes, effects, intentions, reactions	Coverage and ambiguity
ATOMIC 2020	Expanded inferential commonsense	Commonsense knowledge graph with symbolic/neural orientation	Connects structured commonsense and neural knowledge modelling	Scale, quality, and generalization remain open
Knowledge graphs + LLMs	Hybrid world knowledge	Graphs, embeddings, language models	Combines retrieval, structure, and generation	Integration, provenance, and unsupported generation

6.1 ConceptNet

ConceptNet 5.5 represents words and phrases as concepts connected by labelled edges and was designed to provide general knowledge useful for language understanding. The resource draws on multiple sources, including expert-created resources and crowdsourcing, and its graph

structure makes relational commonsense directly accessible to computational systems (Speer et al., 2017). An assertion such as (cup, UsedFor, drinking) demonstrates the basic pattern: a concept, a typed relation, and another concept. This structure supports traversal, semantic association, and integration with distributional representations.



Illustrative structures derived from examples and descriptions in the article; not a data extract.

Figure 3. Illustrative relational and event-centred commonsense structures: ConceptNet-style and ATOMIC-style representations.

6.2 ATOMIC and ATOMIC 2020

ATOMIC shifted attention from taxonomic knowledge toward event-centred inferential knowledge. Its typed if-then relations encode likely causes, effects, intentions, and reactions, thereby addressing dimensions of commonsense that cannot be captured by simple class hierarchies (Sap et al., 2019). ATOMIC 2020 extended this direction and explicitly examined the utility of commonsense knowledge graphs for learning implicit knowledge about new entities and events (Hwang et al., 2021). These resources illustrate the importance of representing not only what things are, but also what tends to happen, what agents may want, and how events affect participants.

7. Commonsense Reasoning in Modern AI

7.1 Natural-language processing

Natural-language understanding frequently

requires background knowledge. Coreference resolution, question answering, dialogue, summarization, and interpretation of implied meaning may depend on information that is absent from the immediate text. Knowledge graphs can provide structured retrieval while neural models provide contextual language interpretation. Their integration can support systems that retrieve explicit relations while using learned models for flexible language understanding.

7.2 Planning and intelligent agents

Planning requires assumptions about actions, preconditions, effects, persistence, affordances, obstacles, containment, and goals. Commonsense therefore provides a bridge between symbolic descriptions of a situation and decisions about what can or should happen next.




Figure 4. Challenge landscape synthesized from the limitations discussed in the review.

7.3 Robotics and embodied AI

Robots operate in physical environments where commonsense expectations can prevent unreasonable actions. Spatial and physical

knowledge must be connected to sensorimotor experience, making grounding a central issue. Abstract textual knowledge is useful only when the system can connect it to actual entities, states, and actions.

7.4 Explainability and evidence

Structured knowledge can make some AI inferences easier to inspect because relevant entities, relations, evidence, and inference paths

can be represented explicitly. However, an apparently clear graph explanation is valuable only if the underlying knowledge is accurate, relevant, current, and appropriately contextualized.

8. Challenges and Limitations

Challenge	Why it matters
Knowledge acquisition	Manual encoding is expensive; automatic extraction can introduce noise, ambiguity, duplication, and context loss.
Coverage	The open-ended number of everyday facts and situations makes complete coverage unrealistic.
Context	A proposition can change relevance depending on time, place, agents, goals, and discourse context.
Inconsistency	Different sources may contain contradictory or incompatible assertions.
Non-monotonicity	New evidence can invalidate an earlier reasonable conclusion.
Temporal change	Social conventions and knowledge can evolve over time.
Cultural variation	Everyday assumptions differ across communities and languages.
Grounding	Textual knowledge may not correspond reliably to physical or social reality.
Scalability	Reasoning across very large graphs can be computationally expensive.
Evaluation	Narrow benchmarks may not measure robust real-world understanding.
Bias and provenance	Crowdsourced and web-derived knowledge may encode social or cultural biases.
Unsupported generation	Neural generators can produce plausible claims without sufficient evidence.

8.1 The knowledge-acquisition bottleneck

The gap between the amounts of commonsense humans possess and the amount that can be explicitly encoded remains fundamental. Manual construction can yield carefully controlled knowledge but is slow and costly. Automatic extraction is scalable but must address ambiguity, source reliability, duplicate facts, and context. Crowdsourcing increases breadth but

introduces variation in interpretation and annotation quality.

8.2 The reasoning bottleneck

Even when relevant knowledge exists, an AI system must identify which facts matter, combine them in an appropriate order, resolve exceptions, and determine whether a conclusion is sufficiently supported. Commonsense is

therefore a reasoning problem as much as a representation problem.


8.3 The evaluation problem

Accuracy on a narrow commonsense benchmark is insufficient evidence of robust everyday reasoning. More informative evaluation should combine benchmark performance with robustness, calibration, contradiction detection, temporal updating, explanation faithfulness, cross-cultural validity, and realistic interactive tasks. Such evaluation is especially important when commonsense systems are embedded in educational or advisory applications.

9. A Proposed Hybrid Framework for Commonsense-Enabled AI

The synthesis in this review supports a layered architecture rather than a single universal representation. The proposed framework contains five interacting layers: (1) input interpretation; (2) an explicit commonsense knowledge layer; (3) neural retrieval and semantic matching; (4) a reasoning layer combining symbolic, defeasible, and probabilistic mechanisms; and (5) an explanation and provenance layer. The framework is conceptual and has not been experimentally validated in this article.

Layer	Core function	Possible technologies	Primary evaluation focus
1. Input/interpretation	Extract entities, events, intent, and context	Transformers, NLP pipelines	Entity/event and context accuracy
2. Knowledge	Store explicit commonsense relations and metadata	Knowledge graph, ontology, rules	Coverage, consistency, provenance
3. Retrieval	Find relevant facts and analogies	Embeddings, vector search, graph traversal	Recall and relevance
4. Reasoning	Derive or revise conclusions	Logic, defeasible rules, probabilistic models, neural inference	Validity, calibration, robustness
5.Explanation/provenance	Expose evidence and inference path	Provenance graph, rule trace, citations	Faithfulness and understandability

Proposed layered hybrid framework for commonsense-enabled AI

Conceptual research design; the article states that the framework has not been experimentally validated.

Figure 5. Proposed layered hybrid framework for commonsense-enabled AI.

9.1 Conceptual model

The framework can be summarized as:

Intelligent Response = f(Context, Retrieved Knowledge, Representation, Reasoning, Evidence, Action Constraints)

This expression is a research-design principle rather than a statistical or mathematical model. It emphasizes that reliable commonsense behaviour depends on interactions among contextual information, retrieved knowledge, representational form, reasoning mechanisms, evidence, and constraints on action.

10. Implications for Educational AI and Nigerian Contexts

Educational AI provides a particularly relevant application domain because students and teachers communicate through incomplete, contextual, and culturally situated language. An

educational assistant may need to interpret an underspecified question, infer a likely misconception, recognize classroom conventions, and select an explanation appropriate to the learner's immediate context. Commonsense representation can support these functions when it is used as contextual evidence rather than as an inflexible set of universal rules.

10.1 Intelligent tutoring and academic advising

A commonsense-aware educational assistant could represent classroom situations, learner intentions, ordinary academic workflows, and common relationships among tasks and outcomes. For example, a system could use procedural and temporal knowledge to interpret references to assignments, deadlines, feedback, and revision, while an explanation layer could show which retrieved knowledge supported a response.

10.2 Nigerian and African educational contexts

A major research opportunity is the development of context-sensitive commonsense resources for Nigerian educational environments. Such resources could represent local institutional practices, multilingual expressions, common classroom situations, culturally grounded examples, and educational workflows. The objective should not be to encode stereotypes as universal facts. Instead, the system should preserve provenance and distinguish context-specific social knowledge from more general physical, logical, or mathematical regularities.

10.3 Potential research applications

- commonsense knowledge graphs for Nigerian educational scenarios;
- multilingual educational commonsense resources covering English and selected Nigerian languages;
- context-aware AI tutors for teacher-education programmes;
- knowledge-grounded academic question-answering systems;
- explainable educational recommendation systems using graph-based evidence; and
- benchmark suites for commonsense reasoning in Nigerian classroom contexts.

11. Research Directions and Opportunities

- Hybrid knowledge architectures that integrate knowledge graphs, symbolic rules, probabilistic inference, and neural language models.
- Context-aware commonsense representations that explicitly encode time, place, task, agent, discourse, and conversational context.
- Continual commonsense learning that updates knowledge while preserving provenance and reducing catastrophic forgetting.
- Uncertainty-aware knowledge graphs that attach confidence, source reliability, and temporal validity to assertions.
- Culturally grounded commonsense resources that support African and Nigerian contexts without assuming that one cultural model is universal.
- Commonsense for educational AI, especially tutoring, explanation, classroom dialogue, academic advising, and adaptive learning.
- Evaluation beyond accuracy, including calibration, consistency, contradiction handling, explanation quality, temporal updating, and out-of-distribution reasoning.
- Grounded commonsense that connects textual knowledge with visual, spatial, sensor, and interaction data.

11.1 Proposed research model

Element	Proposed variables/components
Independent variables	Representation strategy; knowledge source; reasoning strategy; contextual metadata
Intermediate mechanisms	Retrieval relevance; inference depth; uncertainty estimation; contradiction handling
Dependent outcomes	Commonsense accuracy; robustness; calibration; explanation quality; task success
Moderating factors	Language; cultural context; domain; temporal drift; data quality

12. Discussion

The central implication of the reviewed literature is that commonsense should be treated as a systems problem rather than as a property that can be guaranteed by choosing one representation family. Symbolic approaches contribute explicitness, traceability, and formal inference. Frames and scripts capture recurring situations. Ontologies support semantic precision and interoperability. Knowledge graphs provide relational organization and integration. Probabilistic approaches provide mechanisms for uncertainty. Neural models contribute broad learned coverage and flexible language interaction. Neuro-symbolic architectures seek to connect these strengths.

At the same time, combining representations does not automatically solve the problem. A hybrid architecture introduces its own research questions: how knowledge is selected, how conflicting assertions are reconciled, how confidence is calibrated, how temporal validity is maintained, and how cultural assumptions are represented. The proposed framework therefore treats provenance, context, and evaluation as first-class components rather than optional additions.

For doctoral research, the most defensible contribution may be a focused advancement in one layer of this architecture—such as culturally grounded knowledge acquisition, context-aware retrieval, uncertainty-aware graph reasoning, or evaluation in educational environments—followed by empirical validation against clearly defined baselines. This framing turns the broad commonsense problem into a set of testable research questions without assuming that a complete representation of human commonsense is achievable.

13. Limitations of the Review

This article is a conceptual review derived from an existing doctoral seminar rather than a systematic review with a preregistered search protocol, formal inclusion/exclusion criteria, or quantitative synthesis. Consequently, the literature discussed is selective and organized around the themes developed in the source

seminar. The proposed architecture is conceptual and has not been implemented or benchmarked within this study. Before submission to a specific journal, the manuscript should be further aligned with that journal's word limit, reference style, reporting requirements, and expectations for systematic or scoping reviews if applicable.

14. Conclusion

Commonsense knowledge remains a foundational but difficult component of artificial intelligence because intelligent behaviour requires more than literal pattern matching. Systems must interpret incomplete information, reason with defaults and uncertainty, understand causal and social relations, adapt to context, and connect abstract representations with real-world situations.

The historical record shows that each major representation paradigm solves part of the problem. Logic provides explicit formal inference; frames and scripts represent stereotyped situations; semantic networks and ontologies organize concepts and relations; knowledge graphs provide scalable relational structure; probabilistic models represent uncertainty; neural models provide broad learned coverage; and neuro-symbolic systems attempt to integrate complementary strengths. Resources such as Cyc, ConceptNet, ATOMIC, and ATOMIC 2020 demonstrate both the progress made and the continuing problems of coverage, quality, grounding, provenance, and evaluation.

The review consequently proposes a layered architecture that combines contextual interpretation, explicit commonsense knowledge, neural retrieval, structured and probabilistic reasoning, and explanation/provenance. The framework should be regarded as a research agenda rather than a validated solution. For educational AI, especially in Nigerian and African contexts, it offers a basis for developing context-aware tutors, knowledge-grounded academic assistants, culturally responsive resources, and evaluation benchmarks. Future work should operationalize these ideas through well-defined datasets, reproducible systems, and empirical evaluation

across realistic and culturally varied scenarios.

Declarations

Funding: This research received no external funding.

Conflict of interest: Conflict of Interest

The authors declare that they have no conflict of interest regarding the publication of this paper.

Data availability: No original dataset was created or analyzed in this conceptual review.

Author contributions:

Ndiana Okon Asuquo and Obonguko, Ubon Asuquo jointly contributed to the conceptualization, literature review, analysis, and preparation of this manuscript. Both authors participated in reviewing and editing the manuscript and approved the final version for publication.

Ndiana Okon Asuquo write the original draft while Obonguko, Ubon Asuquo was responsible for supervision, methodology, data analysis

References

- ACM/IEEE-CS. (2024). CS2023: Computer Science Curricula—Artificial Intelligence and Knowledge Representation and Reasoning guidance. Association for Computing Machinery and IEEE Computer Society.
- Cycorp. (n.d.). Cyc platform and knowledge base documentation. <https://cyc.com/platform/>
- Davis, E. (1990). Representations of commonsense knowledge. Morgan Kaufmann. <https://doi.org/10.1016/C2013-0-08296-5>

- Davis, E., & Marcus, G. (2015). Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM*, 58(9), 92–103. <https://doi.org/10.1145/2701413>
- Hwang, J. D., Bhagavatula, C., Le Bras, R., Da, J., Sakaguchi, K., Bosselut, A., & Choi, Y. (2021). (Comet-) ATOMIC 2020: On symbolic and neural commonsense knowledge graphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7), 6384–6392. <https://doi.org/10.1609/aaai.v35i7.16792>
- Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2), 494–514. <https://doi.org/10.1109/TNNLS.2021.3070843>
- Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- Sap, M., Le Bras, R., Allaway, E., Bhagavatula, C., Lourie, N., Rashkin, H., Roof, B., Smith, N. A., & Choi, Y. (2019). ATOMIC: An atlas of machine commonsense for if-then reasoning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(1), 3027–3035. <https://doi.org/10.1609/aaai.v33i01.33013027>
- Speer, R., Chin, J., & Havasi, C. (2017). ConceptNet 5.5: An open multilingual graph of general knowledge. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1), 4444–4451. <https://doi.org/10.1609/aaai.v31i1.11164>