From Universal Knowledge Graphs to Contextual Semantic Contracts

Himanshu Beniwal[1 [0000−0001−6389−576X]] and Michael Färber[1 [0000−0001−5458−8645]]

ScaDS.AI / Technische Universität Dresden, Dresden 01069, Germany
{himanshu.beniwal, michael.faerber}@tu-dresden.de

Abstract. The Semantic Web has transformed the Web into a machine-readable ecosystem through ontologies, linked data, and knowledge graphs, enabling structured representation and reasoning over factual knowledge. However, the emergence of foundation models and generative AI has fundamentally changed how knowledge is consumed, generated, and interpreted. Modern AI systems increasingly fail not because they lack factual knowledge, but because they struggle to interpret that knowledge under diverse temporal, cultural, legal, and social contexts. We argue that this limitation stems from a deeper representational gap: the Semantic Web primarily models facts, while future AI requires explicit representations of contextual meaning. We envision the next generation of semantic technologies as evolving from a Web of Facts to a Web of Context, where context becomes a first-class semantic object that governs interpretation rather than merely annotating it. We outline a research agenda for context-centric semantic representations and argue that representing meaning, rather than facts alone, should become the defining challenge for the next twenty-five years of Semantic Web research.

Keywords: Semantic Web · Web of Context · Semantic Interoperability.

1 Introduction

Twenty-five years ago, the Semantic Web introduced a transformative vision: to make the Web understandable by machines through structured representations of knowledge [9, 7]. Ontologies, RDF, OWL, Linked Open Data, and Knowledge Graphs collectively shifted the Web from a collection of documents to a network of machine-readable facts [3]. This vision has profoundly influenced search engines, biomedical informatics, digital libraries, enterprise systems, and knowledge-intensive AI applications [8, 4, 5].

Today, however, the landscape has fundamentally changed. Large Language Models (LLMs) no longer consume isolated triples or curated ontologies; instead, they integrate information from billions of heterogeneous sources while generating language, reasoning over incomplete evidence, and synthesizing new knowledge. The challenge is therefore no longer whether machines can access facts, but whether they can interpret those facts appropriately under different contexts.



Fig. 1. Illustrative examples of culturally aware image transformations across Hindu and Muslim contexts. The four quadrants show: (top left) transforming a depiction of Lord Ganesh on an Indian carpet; (bottom left) modifying religious imagery on commercial footwear; (top right) addressing the stereotype of a Muslim person consuming alcohol; and (bottom right) mitigating imagery of a Hindu priest consuming beef. ***Takeaway***: *multimodal systems can reproduce harmful religious stereotypes through inappropriate object-deity associations and misrepresentations of religious practice.*

This distinction is subtle but fundamental. The Semantic Web has primarily focused on representing what is true [8], whereas modern AI increasingly requires understanding when, for whom, under which assumptions, and in which cultural or temporal context that truth should be interpreted. Related concerns are partly addressed by principles for metadata, provenance, and data reuse [12], which do not target cultural or social interpretation.

1.1 Motivation

Consider the following example in Figure 1: A knowledge graph may correctly represent the fact: Ganesh → instance of → Hindu deity, and another fact, Shoes → instance of → commercial product, both statements are factually correct.

Now, imagine a generative model producing an advertisement depicting Lord Ganesh printed on footwear. The underlying knowledge is not incorrect; the generated image is visually coherent and factually consistent. Nevertheless, the content is deeply offensive within many Hindu communities because the relationship between the entities carries cultural meaning that is absent from existing semantic representations.

Similarly, a multimodal model may generate images depicting a Muslim individual consuming alcohol or a Hindu priest eating beef. These outputs do not necessarily violate factual knowledge; rather, they violate community-specific cultural and religious norms. These examples illustrate that multimodal systems can perpetuate harmful stereotypes and culturally insensitive associations even when the underlying entities are correctly recognized [1, 11].

These examples reveal an important limitation of today's Semantic Web. Existing semantic technologies are highly effective at representing entities, relations, and logical constraints, yet they rarely encode contextual meaning, including

cultural appropriateness, temporal validity, social norms, provenance, or acceptable usage. Consequently, these forms of knowledge remain implicit within neural model parameters rather than being explicitly represented, shared, or verified.

1.2 The Emerging Gap

This observation exposes a growing mismatch between the assumptions of the traditional Semantic Web and the needs of modern AI [2, 6]. Knowledge Graphs answer questions such as: ***(1)*** What is this entity? ***(2)*** How are two entities related? ***(3)*** Which logical facts can be inferred?

Modern AI systems, however, increasingly require answers to additional questions: ***(1)*** Is this interpretation appropriate in the current cultural context, and does this fact remain valid today? ***(2)*** Under which legal, religious, or social setting does this relation hold? ***(3)*** Are multiple interpretations simultaneously acceptable? and ***(4)*** How should conflicting viewpoints be represented rather than eliminated?

These questions are fundamentally semantic, yet current semantic representations provide limited support for expressing them.

1.3 Our Vision

We argue that the next evolution of the Semantic Web should not primarily focus on representing more knowledge, but on representing richer meaning.

Instead of treating semantics as a collection of universally valid facts, we envision semantic representations that explicitly capture context, norms, provenance, temporal validity, uncertainty, and community-specific interpretations alongside factual knowledge. Such representations would enable AI systems to reason not only about what exists, but also about how information should be interpreted under different circumstances.

In other words, we propose a shift, "*from a Web of Facts to a Web of Context*". This shift does not replace existing Semantic Web technologies; rather, it extends them with a complementary semantic layer that models meaning beyond factual correctness.

1.4 Research Questions

This vision raises several fundamental questions that, in our opinion, should define the next generation of Semantic Web research.

- RQ1. Can semantic representations explicitly model contextual meaning including cultural, temporal, legal, and social dimensions—rather than only factual knowledge?
- RQ2. What new semantic abstractions are required to represent norms, acceptable interpretations, provenance, and uncertainty as first-class semantic objects?

- RQ3. Should future semantic languages evolve beyond representing static relationships between entities toward representing context-dependent interpretations of those relationships?
- RQ4. How should semantic technologies represent disagreement, plurality, and community-specific viewpoints without forcing a single universal truth?
- RQ5. What new benchmarks and evaluation paradigms are needed to assess semantic representations based on contextual understanding rather than factual completeness alone?

2 The Semantic Gap: When Facts Are Not Enough

The Semantic Web was envisioned as a means for knowledge integration and retrieval; it primarily models facts rather than their interpretations. This paradigm assumes that semantics can be adequately captured through entities, relationships, and logical constraints. This assumption has proven highly effective for integration and retrieval, but it leaves the conditions of interpretation largely outside the representation itself.

The emergence of foundation models has fundamentally changed this landscape. Modern AI systems no longer retrieve isolated facts; they synthesize information, generate content, and reason across diverse contexts. Consequently, the challenge has shifted from retrieving facts to interpreting them appropriately.

Consider the statement, "*A lion is a dangerous animal*". The fact remains unchanged, yet its meaning depends on context—it may describe ecological behavior, educational knowledge, metaphorical language, or operational risk. Likewise, an image depicting a religious figure on commercial products may be factually accurate but culturally inappropriate. These examples illustrate that factual correctness alone is insufficient for semantic understanding.

This exposes a fundamental semantic gap: current semantic technologies effectively represent what is true but provide limited support for representing when, where, for whom, and under which conditions that knowledge should be interpreted. We argue that addressing this gap is the next frontier of the Semantic Web [1].

- Temporal context: Medical recommendations change as clinical evidence evolves, necessitating clear distinctions between current guidance and historical knowledge.
- Legal context: Regulations differ across jurisdictions, making identical actions permissible in one country and prohibited in another.
- Scientific context: Multiple competing hypotheses may coexist until new evidence emerges.
- Social context: Community norms, etiquette, and acceptable behavior often differ across cultures despite referring to the same underlying entities.

Current semantic technologies can often represent these aspects indirectly through additional properties or metadata. However, they generally treat context as supplementary information rather than as a first-class semantic object

that directly influences interpretation and reasoning. As a result, contextual knowledge remains fragmented across application-specific schemas or is implicitly encoded within neural model parameters, making it difficult to share, verify, or reason over consistently.

This limitation becomes increasingly significant as AI systems move beyond information retrieval toward content generation, decision support, and collaborative reasoning. In these settings, correctness is no longer determined solely by whether a statement is factually accurate, but also by whether it is appropriate, applicable, timely, and interpretable within the intended context.

We therefore argue that the next generation of semantic technologies should expand the notion of semantics itself. Rather than representing only entities and their relationships, semantic representations should explicitly capture the contextual conditions under which knowledge is valid, applicable, and meaningful. This shift—from representing facts to representing contextual meaning—forms the conceptual foundation of our vision.

3 Reimagining the Semantic Web: From Knowledge Representation to Meaning Representation

The next generation of the Semantic Web should evolve from representing what is true to representing how truth should be interpreted.

For over two decades, semantic technologies have successfully represented entities, relationships, and logical facts through ontologies, linked data, and knowledge graphs. While these representations enable interoperability and reasoning, they assume that semantics is largely independent of context [10].

The era of generative AI challenges this assumption. Modern AI systems must determine not only whether information is correct, but also whether it is appropriate, trustworthy, temporally valid, culturally acceptable, and relevant. These properties emerge from context, not from facts alone.

We therefore argue that context should become a first-class semantic object. Semantic representations should explicitly capture the conditions under which knowledge is interpreted, including temporal validity, cultural norms, legal jurisdictions, provenance, uncertainty, and intended audience.

This shift complements—not replaces—existing semantic technologies. Knowledge graphs describe what; contextual semantics explains when, where, for whom, and under which assumptions that knowledge applies. Consequently, semantic interoperability should extend beyond sharing facts to sharing the contextual conditions that govern their interpretation.

Ultimately, we envision the next generation of the Semantic Web as an infrastructure for meaning representation, enabling AI systems to reason not only over facts, but over their context.

3.1 Why this constitutes a paradigm shift

The first generation of the Semantic Web asked: "How can machines understand data?", the next generation should ask: "How can machines understand mean-

ing?". This distinction is subtle but transformative. Knowledge-graph assertions are often modeled as relatively static and context-independent, whereas meaning is dynamic, contextual, and often pluralistic. A Semantic Web that represents meaning would enable AI systems to reason over context, negotiate multiple valid interpretations, and adapt knowledge to diverse cultural, temporal, legal, and social settings without assuming that a single representation is universally sufficient.

4 A Research Agenda for the Next Generation of the Semantic Web

The transition from a Web of Facts to a Web of Context is not merely an extension of existing semantic technologies; it represents a fundamental shift in what semantic systems are expected to capture. Realizing this vision requires rethinking several long-standing assumptions that have guided Semantic Web research for more than two decades. Rather than proposing a single framework or ontology, we identify a set of foundational research directions that we believe should define the next generation of semantic technologies.

4.1 Context as a First-Class Semantic Object

Current semantic representations primarily describe entities and relationships. Context is often represented indirectly through metadata or application-specific annotations, limiting its role in semantic reasoning.

Future semantic models should treat context itself as a first-class semantic object. Context should be explicitly represented, queried, reasoned over, and shared across systems. Such context may include temporal validity, geographical scope, legal jurisdiction, cultural practices, intended audience, environmental conditions, or application-specific assumptions. By elevating context to the same level as entities and relations, semantic systems can move beyond universally interpreted facts toward context-aware knowledge representations.

4.2 Representing Multiple Valid Interpretations

Many mainstream semantic representations do not explicitly encode multiple context-dependent interpretations of the same statement. However, many real-world concepts admit multiple valid interpretations that depend on community, discipline, culture, or historical period.

Future semantic technologies should enable multiple contextual interpretations to coexist without forcing a single universal representation. Rather than treating disagreement as inconsistency, semantic systems should explicitly model competing viewpoints together with the conditions under which each interpretation applies.

Such representations would allow AI systems to acknowledge diversity instead of collapsing complex social knowledge into simplified universal statements.

4.3 Semantics Beyond Truth

Semantic Web research has historically focused on representing factual correctness and logical consistency. While these remain essential, future AI systems increasingly require additional semantic dimensions.

Questions such as ***(1)*** Is this recommendation appropriate? ***(2)*** Is this information still applicable today? ***(3)*** Should this knowledge be presented to this audience? and ***(4)*** Does this interpretation respect community norms? cannot be answered through factual reasoning alone.

We therefore argue that future semantic representations should incorporate dimensions such as applicability, appropriateness, provenance, uncertainty, temporal validity, and normative constraints as integral components of semantic meaning rather than auxiliary metadata.

4.4 Measuring Semantic Understanding

The evaluation of semantic technologies has traditionally emphasized logical consistency, ontology quality, reasoning performance, and query efficiency. These metrics remain important but are insufficient for evaluating contextual understanding.

Future benchmarks should assess whether semantic systems can correctly interpret knowledge across varying temporal, cultural, legal, and social contexts. Evaluation should move beyond verifying whether a system retrieves the correct fact toward assessing whether it produces the most appropriate interpretation under a given context.

This shift will require entirely new benchmark datasets, evaluation protocols, and community-wide shared tasks that reflect the contextual nature of meaning.

5 Discussion

The ideas presented in this paper intentionally avoid proposing a new ontology language, reasoning engine, or knowledge graph architecture. Instead, they advocate a conceptual reorientation of the Semantic Web itself. The Web has become increasingly successful at representing what exists; the next challenge is to represent what information means under different circumstances.

This distinction may appear incremental at first glance, yet it fundamentally changes the role of semantic technologies in future AI ecosystems. Rather than serving primarily as repositories of structured knowledge, semantic systems become infrastructures for representing contextual meaning, enabling AI systems to interpret, adapt, and communicate knowledge responsibly across diverse domains and communities.

6 Conclusion

Twenty-five years ago, the Semantic Web taught machines what the world knows. The next twenty-five years should enable machines to understand what that

knowledge means. We believe that this transition—from a Web of Facts to a Web of Context—represents the next fundamental frontier for semantic technologies and offers a roadmap for reimagining the Semantic Web in the age of generative AI.

Acknowledgments. This research was funded by the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) and the German Federal Ministry of Research, Technology and Space (BMFTR) through the Software Campus project (01IS23070).

References

1. Adilazuarda, M.F., Mukherjee, S., Lavania, P., Singh, S.S., Aji, A.F., O'Neill, J., Modi, A., Choudhury, M.: Towards measuring and modeling "culture" in LLMs: A survey. In: Al-Onaizan, Y., Bansal, M., Chen, Y.N. (eds.) Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing. pp. 15763–15784. Association for Computational Linguistics, Miami, Florida, USA (Nov 2024). https://doi.org/10.18653/v1/2024.emnlp-main.882, https://aclanthology.org/2024.emnlp-main.882/
2. Antoniou, G., Harmelen, F.V.: A semantic web primer, 2nd edn, cooperative information systems (2008)
3. Bizer, C., Heath, T., Berners-Lee, T.: Linked data-the story so far. In: Linking the world's information: Essays on tim berners-lee's invention of the world wide web, pp. 115–143 (2023)
4. Carroll, J.J., Bizer, C., Hayes, P., Stickler, P.: Named graphs. Journal of Web Semantics **3**(4), 247–267 (2005)
5. Dey, A.K.: Understanding and using context. Personal and ubiquitous computing **5**(1), 4–7 (2001)
6. Fensel, D., Şimşek, U., Angele, K., Huaman, E., Kärle, E., Panasiuk, O., Toma, I., Umbrich, J., Wahler, A.: Introduction: what is a knowledge graph? In: Knowledge graphs: Methodology, tools and selected use cases, pp. 1–10. Springer (2020)
7. Gruber, T.R.: A translation approach to portable ontology specifications. Knowledge acquisition **5**(2), 199–220 (1993)
8. Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G.D., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., et al.: Knowledge graphs. ACM Computing Surveys (Csur) **54**(4), 1–37 (2021)
9. Lassila, O., Hendler, J., Berners-Lee, T.: The semantic web. Scientific american **284**(5), 34–43 (2001)
10. Lebo, T., Sahoo, S., McGuinness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J.: Prov-o: The prov ontology (2013)
11. Schneider, F., Sitaram, S.: M5–a diverse benchmark to assess the performance of large multimodal models across multilingual and multicultural vision-language tasks. In: Findings of the Association for Computational Linguistics: EMNLP 2024. pp. 4309–4345 (2024)
12. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.W., da Silva Santos, L.B., Bourne, P.E., et al.: The fair guiding principles for scientific data management and stewardship. Scientific data **3**(1), 160018 (2016)