

HTTP for Meaning: Towards Linguistically Complete Knowledge Representations with KNOWLEDGE STATE PROTOCOL

Ian Yuen

Independent Researcher
jardinianr@gmail.com

Abstract

The choice of knowledge representation (KR) determines what a knowledge base (KB) can store, and how it can be retrieved, and reasoned over. Recent Knowledge-Graph (KG) Retrieval-Augmented Generation (RAG) systems show that *explicit symbolic structure* can complement the implicit knowledge encoded in LLMs. However, KGs capture only part of a text’s linguistic meaning and often leave **conceptual** and **temporal** context implicit, leading to recurrent reasoning failures. Existing approaches address these gaps through task-specific extensions, but do not provide a unified framework for coordinating conceptual, instance-level, and temporal structure. We envision the KNOWLEDGE STATE PROTOCOL (KSP), which unifies canonical KRs to make the known linguistic dimensions of a text explicit. By standardizing the interoperation between KRs, KSP provides a common, extensible, and open foundation for *structured* agent memory and neuro-symbolic training across modalities and languages.

1 Introduction

Natural language encodes multiple layers of meaning, whereas any single knowledge representation (KR) externalizes only part of that structure (Table 1). Ontologies make domain concepts, taxonomies, and logical constraints explicit; KGs represent facts stated in source text; and Abstract Meaning Representation (AMR) models sentence-level events and the semantic relations between participants (Banarescu et al., 2013). Each representation has independently supported question answering (QA) systems (Lopez et al., 2005; Berant et al., 2013; Kapanipathi et al., 2021).

However, no single KR captures all *linguistic dimensions* (Table 1), as in the case of KG-based RAG. Although it improves on text-based RAG by representing entities and factual relations symbolically (Edge et al., 2024), it still leaves two recurring

gaps: **conceptual context** (Zhong et al., 2024) and **temporal and event context** (Cai et al., 2023). KGs encode object-relation-object triples which can capture the explicitly stated facts in text, but not always the implied conceptual (e.g., "X is a type of Y") and temporal context (e.g., "A occurred before B") with which humans interpret text.

The limitations of relying on partial or implicit representations are especially consequential in LLM-based systems. Although LLMs achieve strong performance on QA benchmarks (OpenAI, 2026), their internal knowledge is difficult to inspect, attribute, constrain, or update (Lewis et al., 2020). Moreover, in enterprise and specialized settings, the required knowledge may be private, domain-specific, or continuously changing and therefore unavailable during model training (Kandpal et al., 2023; Mallen et al., 2023). Addressing the **conceptual** and **temporal context** gaps in KG-RAG could therefore improve QA performance in LLM systems.

Recent systems address these omissions by augmenting KGs with specialized structures for additional context. Temporal QA methods add timestamps, intervals, dynamic updates, or event-centric relations to preserve how facts evolve (Li et al., 2025; Zhang et al., 2026; Sun et al., 2025; Han et al., 2025; Wang et al., 2026); Conceptual QA methods augment graph-based retrieval and reasoning with ontology-derived concepts, relations, and schemas (Liu et al., 2025; Sharma et al., 2025). Other systems go further by combining entity, concept, event, and temporal structures within a bespoke domain-specific graph or meaning representation (de Martim, 2025; Bai et al., 2026).

Although effective, these approaches add missing capabilities separately or re-encode conceptual, factual, and event-temporal distinctions within a single custom graph schema, rather than coordinating the established KRs that already formalize them. Ontologies can contextualize data as

Linguistic capacity in natural language	KG representation	AMR representation	Ontology representation	Best-fit KR
Factual reference: language refers to entities, relations, and states of affairs.	Strong: represents entity–relation facts.	Moderate: represents sentence-level concepts and roles.	Moderate: defines conceptual types and relations.	KG is strongest for durable factual memory, but not full sentence meaning or abstraction.
Compositional event meaning: predicates, arguments, modifiers, and scope combine into structured meanings.	Weak to moderate: triples often flatten nested or scoped meaning.	Strong: explicitly represents predicate–argument and event structure.	Moderate: provides reusable conceptual primitives, but not full sentential structure.	AMR captures local semantic structure that KG and ontology often miss.
Conceptual abstraction: language expresses classes, roles, norms, constraints, and generalizations.	Moderate: can attach types, but schemas are often incomplete.	Weak to moderate: can express abstract predicates, but not a domain theory.	Strong: represents classes, properties, constraints, inheritance, and equivalence.	Ontology supplies the conceptual layer missing from instance-oriented graphs.
Displacement and temporality: language refers to past, future, absent, hypothetical, or changing states.	Weak unless time and state are explicitly modeled.	Moderate to strong: captures event-level temporal cues, but needs coordination.	Moderate: can define states, events, intervals, and temporal relations.	AMRs structured as a ledger can provide both event extraction and explicit temporal modeling.
Modality, uncertainty, and falsehood: language expresses beliefs, claims, hypotheticals, negation, fiction, and error.	Weak: plain triples risk converting claims into facts.	Moderate to strong: naturally represents negation, attribution, and selected modal meanings.	Strong if designed: can model Claim, Belief, Evidence, Source, and Hypothesis.	Requires AMR -style cues plus ontology -style claim modeling.
Reflexiveness and provenance: language can talk about statements, sources, and representations.	Moderate: possible through reification, named graphs, or metadata.	Moderate: can represent reporting and saying events.	Strong: can define source, evidence, statement, and provenance classes.	Ontology enables provenance as a first-class concept rather than metadata.
Productivity: language can express novel meanings beyond a fixed inventory.	Strong: easily adds new entities and relations, though semantics may be underspecified.	Moderate: composes novel event structures but bounded by frames and relation inventory.	Strong but costly: adds precise concepts and relations through governed schema evolution.	KGs maximize extensibility, while ontologies provide semantic control.

Table 1: For each *linguistic dimension* (Hockett, 1960), a different KR is best suited to represent it.

instances of domain concepts and support abstract multi-hop reasoning through logical hierarchies and constraints. AMR graphs can isolate relationships between entities for each sentence event, while an AMR event ledger can model how entity relationships change over time. What is missing is a shared framework that coordinates these complementary structures through a common integration protocol.

2 Knowledge State Protocol (KSP)

We envision KSP to be a common protocol that captures the known *linguistic dimensions* of a text by coordinating explicit symbolic KRs. Such coordination is feasible: prior work has translated AMR into ontology-aligned KGs and used ontologies to guide KG construction from text (Gangemi et al., 2026; van Cauter and Yakovets, 2024).

2.1 Coordinating KRs

Because different KRs make different *linguistic dimensions* of meaning explicit, KSP operationalizes “linguistic completeness” by coordinating representations derived from source text across three layers: (1) a conceptual layer for domain concepts and constraints, (2) an instance layer for entities and facts, and (3) a temporal layer for events and state

changes (Fig. 1). This organization parallels cognitive and text-based agent-memory architectures that distinguish semantic, episodic, and procedural knowledge (Tulving, 1972; Sumers et al., 2023). As the choice of representation shapes both the form and interpretation of knowledge (McLuh, 1964), KSP adopts ontology, KG, and AMR as representative KRs for these layers while allowing KR form to be extended or replaced as representation research evolves.

2.2 Aligning Concepts and Instances

KSP is representation-agnostic: it aligns KRs to support “linguistic completeness” without prescribing their internal structure. Like HTTP, which standardizes communication without dictating how participating systems store their data (Fielding et al., 2022), KSP defines shared primitives and mappings through which heterogeneous KRs can interoperate. Cross-layer traversal relies on one such *shared primitive* common to symbolic KRs — an object or entity represented as a node.

This alignment enables traversal from ontology concept nodes to their corresponding KG instance nodes. In conceptual QA, the queried concept may not appear explicitly in the KG or source text. As shown in Fig. 1, KSP can first map the query to




Figure 1: At time Z , each AMR ledger event is sequentially applied to the KG layer. Instance nodes in the KG layer are linked to concept nodes in the ontology layer. *Note.* AMR is visualized as an entity graph to illustrate KSP *shared primitive* alignment; predicates and relations are represented as directed edges.

Concept Y, traverse the ontology, and then identify Instance A in the KG to provide an answer.

2.3 Modeling Temporal Changes

KSP makes the sequence of events and their effects on entities explicit through an event ledger. The ledger serves as the version-control entry point, indexing the KG state associated with each point in time and supporting provenance and reasoning over how knowledge evolves.

Each sentence-level event is parsed into AMR, whose predicate–argument structure captures what happened, who participated, and under what conditions, including distinctions such as negation and modal predicates. KSP anchors AMR event participants to corresponding KG entity nodes and translates the event predicate into a graph operation that updates those entities and their relations.

For a temporal query at a specific time interval ($\text{Time}=Z$ in Fig. 1), the relevant operations for each event can be applied forward or reversed to reconstruct the corresponding KG state. Each reconstructed KG therefore represents the knowledge state at a particular point in time.

2.4 Aligning KSP Instances

A KSP instance is a concrete deployment of the protocol, with its own KR realizations, schemas, and data. To support interoperability, each instance implements KSP’s shared primitives and exposes explicit, versioned schemas. KSP uses an align-




Figure 2: Palantir Foundry as an analogue to KSP. Foundry uses Ontology as a semantic index for multimodal data sources such as KG. Change over time is managed by Version Control.

ment module to construct versioned *schema mappings* between instances. Candidate mappings are checked against ontology constraints and, when uncertain, reviewed by humans before publication. Existing techniques can already propose many correspondences automatically, although fully reliable alignment remains an open problem (Lushnei et al., 2026). The OBO Foundry provides practical precedent for governed interoperability among independently maintained ontologies (Smith et al., 2007). Because KSP is alignment-method agnostic, improved techniques can be adopted without changing the protocol.

3 Analogue: Palantir Foundry

Palantir Foundry provides a production analogue at scale for KSP’s central premise: knowledge becomes more useful for retrieval, reasoning, and action when semantic structure implicit in raw data is made explicit (Fig. 2). Foundry’s Ontology maps heterogeneous enterprise data to a defined concept object space and vocabulary, creating a shared symbolic semantic layer used by humans, agents and systems (Palantir Technologies Inc., 2022, 2026c). This parallels KSP’s use of Ontology as a concept layer for instances in KGs.

Foundry also illustrates the value of representing state changes explicitly. Version Control and Branching make operational updates and system evolution inspectable and controlled (Palantir Technologies Inc., 2026a,b). This is analogous to KSP’s AMR event ledger.

KSP generalizes these principles into an open,

representation-independent protocol in which conceptual, instance, and event representations can be linked, inspected, and extended across KBs.

4 Challenges

4.1 KSP Benchmark and Metrics

KSP requires a benchmark of private and dynamic data that jointly tests conceptual, instance-level, and temporal reasoning while localizing errors by layer. We propose *Full-Chain Accuracy*, which counts a prediction as correct only if the system selects the appropriate knowledge state, follows a valid ontology-constrained reasoning path, returns the correct answer, and provides a valid rule-and-evidence derivation. This is feasible because prior work has evaluated intermediate reasoning through explicit temporal evidence chains and verifiable KG paths (Liu et al., 2026; Sui et al., 2025).

4.2 KSP Standardization and Adoption

KSP requires community agreement on shared primitives, interfaces, and governance, while ontology and dataset construction impose upfront costs. However, it can build on existing infrastructure: the OBO Foundry provides openly maintained ontologies and shared development principles, while the Knowledge Graph Change Language supports standardized, version-controlled ontology updates (Jackson et al., 2021; Hegde et al., 2025). Recent LLM-based methods may further reduce adoption costs by supporting the scalable construction of domain knowledge graphs with low human oversight (Chen et al., 2024). These investments are justified if standardization prevents each system from reconstructing the same semantic connections independently. Like HTTP, KSP reduces coordination costs with upfront standardization—KSP can enable knowledge to move reliably across systems.

4.3 KSP Scalability and Implementation

KSP’s event ledger enables past states to be reconstructed by replaying recorded events, but rebuilding derived state from a long history can be costly (Erb et al., 2016; Overeem et al., 2021). Checkpointing can reduce this cost by limiting reconstruction to events recorded after the latest cached state (Burckhardt et al., 2022). Palantir Foundry likewise provides production precedent for maintaining current ontology states alongside action histories at scale (Palantir Technologies, 2026; Palantir Technologies Inc., 2026a). We therefore treat scalable

AMR-to-KG reconstruction as an implementation challenge; efficient event translation and temporal retrieval remain open research questions.

5 Research Directions

We envision KSP to make representation of knowledge complete, extensible and modular. KSP can address the failure modes of current KG-RAG and agent-memory systems that lack *structured temporal* and *conceptual* grounding.

Neuro-symbolic Training: KSP suggests a direction for neuro-symbolic training by turning aligned text and knowledge representations into ordered state transitions. DRAGON and KG-BiLM provide precedents for jointly learning from text and graph structure (Yasunaga et al., 2022; Chen et al., 2026). By representing AMR events as graph updates, KSP could extend link-prediction objectives to forecasting future events, the edges they modify, and the resulting knowledge state. Temporal KG models already support future-fact prediction and edge-transition modeling from graph histories (Jin et al., 2020; Han et al., 2021). KSP could further use ontology-level constraints to guide which predicted transitions are semantically valid.

Multi-Modality and Cross-Lingual Representation: Because the *shared primitive* of an object or entity as a node exist across modalities (i.e. scene graphs), KSP can align visual and textual entities and events within shared knowledge states (Ma et al., 2022). The same primitives can connect meaning across languages without erasing language-specific distinctions (Wein and Schneider, 2024). By coordinating these representations, KSP preserves semantic connections across modalities, languages, and changing knowledge, providing a richer foundation for retrieval and reasoning.

6 Conclusion

We identify conceptual and temporal context as persistent gaps in KG-based retrieval and show that existing approaches address them through fragmented, specialized structures. KSP instead coordinates complementary knowledge representations to make these dimensions explicit within a shared knowledge state. This provides a foundation for richer, more interpretable reasoning and for extending structured knowledge across modalities and languages.

7 Limitations

KSP coordinates complementary symbolic KRs to make multiple dimensions of textual meaning explicit, but its performance depends on the accuracy, latency, and cost of the components that construct and align each layer. AMR parsing, text-to-KG construction, and ontology learning have advanced substantially (Bai et al., 2022; Zhang and Soh, 2024; Al-Aswadi et al., 2020), but remain imperfect. Because KSP combines multiple processing stages, errors may also propagate across layers (Bouziani et al., 2024).

The design and optimization of these components fall outside the scope of the protocol definition. Instead, KSP specifies interfaces between KRs without prescribing particular representations or parsers, allowing improved components to be adopted without changing the protocol. Large-scale KG deployments and collaborative ontology ecosystems demonstrate the practical viability of the underlying representations (Noy et al., 2019; Jackson et al., 2021), while AMR has supported a broad range of downstream applications (Wein and Opitz, 2024). Confidence propagation, cross-layer validation, and adaptation to noisy real-world data nevertheless remain open implementation challenges.

References

- Fatima N. Al-Aswadi, Huah Yong Chan, and Keng Hoon Gan. 2020. [Automatic ontology construction from text: A review from shallow to deep learning trend](#). *Artificial Intelligence Review*, 53:3901–3928.
- Jiixin Bai, Wei Fan, Qi Hu, Qing Zong, Chunyang Li, Hong Ting Tsang, Hongyu Luo, Yauwai Yim, Haoyu Huang, Xiao Zhou, Feng Qin, Tianshi Zheng, Xi Peng, Xin Yao, Huiwen Yang, Leijie Wu, Ji Yi, Gong Zhang, Renhai Chen, and Yangqiu Song. 2026. [AutoSchemaKG: Autonomous knowledge graph construction through dynamic schema induction from web-scale corpora](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 20557–20584, San Diego, California, United States. Association for Computational Linguistics.
- Xuefeng Bai, Yulong Chen, and Yue Zhang. 2022. [Graph pre-training for AMR parsing and generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6001–6015, Dublin, Ireland. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Nacime Bouziani, Shubhi Tyagi, Joseph Fisher, Jens Lehmann, and Andrea Pierleoni. 2024. [REXEL: An end-to-end model for document-level relation extraction and entity linking](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*, pages 119–130, Mexico City, Mexico. Association for Computational Linguistics.
- Sebastian Burckhardt, Badrish Chandramouli, Chris Gillum, David Justo, Konstantinos Kallas, Connor McMahon, Christopher S. Meiklejohn, and Xi-angfeng Zhu. 2022. [Netherite: Efficient execution of serverless workflows](#). *Proceedings of the VLDB Endowment*, 15(8):1591–1604.
- Borui Cai, Yong Xiang, Longxiang Gao, He Zhang, Yunfeng Li, and Jianxin Li. 2023. [Temporal knowledge graph completion: A survey](#). In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 6545–6553.
- Hanzhu Chen, Xu Shen, Qitan Lv, Jie Wang, Xiaoqi Ni, and Jieping Ye. 2024. [SAC-KG: Exploiting large language models as skilled automatic constructors for domain knowledge graph](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4345–4360, Bangkok, Thailand. Association for Computational Linguistics.
- Zirui Chen, Xin Wang, Zhao Li, Wenbin Guo, Dongxiao He, Yanbing Li, and Wushour Silamu. 2026. [KG-BiLM: Knowledge graph embedding via bidirectional language models](#). In *Proceedings of the ACM Web Conference 2026*, pages 4290–4301, New York, NY, USA. Association for Computing Machinery.
- Hudson de Martim. 2025. [An ontology-driven graph RAG for legal norms: A hierarchical, temporal, and deterministic approach](#). In *Legal Knowledge and Information Systems – JURIX 2025: The Thirty-Eighth Annual Conference*, volume 416 of *Frontiers in Artificial Intelligence and Applications*, pages 282–287. IOS Press.

- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. [From local to global: A graph RAG approach to query-focused summarization](#). *arXiv preprint arXiv:2404.16130*.
- Benjamin Erb, Gerhard Habiger, and Franz J. Hauck. 2016. [On the potential of event sourcing for retroactive actor-based programming](#). In *Proceedings of the First Workshop on Programming Models and Languages for Distributed Computing*, PMLDC '16, New York, NY, USA. Association for Computing Machinery.
- Roy T. Fielding, Mark Nottingham, and Julian Reschke. 2022. [HTTP semantics](#). RFC 9110, Internet Engineering Task Force.
- Aldo Gangemi, Arianna Graciotti, Antonello Meloni, Andrea Giovanni Nuzzolese, Valentina Presutti, Diego Reforgiato Recupero, and Alessandro Russo. 2026. [Text2AMR2FRED, converting text into RDF/OWL knowledge graphs via abstract meaning representation](#). *Knowledge and Information Systems*, 68:47.
- Jiale Han, Austin Cheung, Yubai Wei, Zheng Yu, Xusheng Wang, Bing Zhu, and Yi Yang. 2025. [RAG meets temporal graphs: Time-sensitive modeling and retrieval for evolving knowledge](#). *Preprint*, arXiv:2510.13590.
- Zhen Han, Zifeng Ding, Yunpu Ma, Yujia Gu, and Volker Tresp. 2021. [Learning neural ordinary equations for forecasting future links on temporal knowledge graphs](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8352–8364, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Harshad Hegde, Jennifer Vendetti, Damien Goutte-Gattat, J. Harry Caulfield, John B. Graybeal, Nomi L. Harris, Naouel Karam, Christian Kindermann, Nicolas Matentzoglou, James A. Overton, Mark A. Musen, and Christopher J. Mungall. 2025. [A change language for ontologies and knowledge graphs](#). *Database*, 2025:baae133.
- Charles F. Hockett. 1960. [The origin of speech](#). *Scientific American*, 203(3):88–97.
- Rebecca Jackson, Nicolas Matentzoglou, James A. Overton, Randi Vita, James P. Balhoff, Pier Luigi Buttigieg, Seth Carbon, Melanie Courtot, Alexander D. Diehl, Damion M. Dooley, William D. Duncan, Nomi L. Harris, Melissa A. Haendel, Suzanna E. Lewis, Darren A. Natale, David Osumi-Sutherland, Alan Ruttenberg, Lynn M. Schriml, Barry Smith, and 6 others. 2021. [OBO foundry in 2021: Operationalizing open data principles to evaluate ontologies](#). *Database*, 2021:baab069.
- Woojeong Jin, Meng Qu, Xisen Jin, and Xiang Ren. 2020. [Recurrent event network: Autoregressive structure inference over temporal knowledge graphs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6669–6683, Online. Association for Computational Linguistics.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. Large language models struggle to learn long-tail knowledge. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 15696–15707. PMLR.
- Pavan Kapanipathi, Ibrahim Abdelaziz, Srinivas Ravishankar, Salim Roukos, Alexander Gray, Ramón Fernández Astudillo, Maria Chang, Cristina Cornelio, Saswati Dana, Achille Fokoue, Dinesh Garg, Alfio Gliozzo, Sairam Gurajada, Hima Karanam, Naweel Khan, Dinesh Khandelwal, Young-Suk Lee, Yunyao Li, Francois Luus, and 11 others. 2021. [Leveraging Abstract Meaning Representation for knowledge base question answering](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3884–3894, Online. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Dong Li, Yichen Niu, Ying Ai, Xiang Zou, Biquing Qi, and Jianxing Liu. 2025. [T-GRAG: A dynamic GraphRAG framework for resolving temporal conflicts and redundancy in knowledge retrieval](#). *Preprint*, arXiv:2508.01680.
- Runxuan Liu, Bei Luo, Jiaqi Li, Baoxin Wang, Ming Liu, Dayong Wu, Shijin Wang, and Bing Qin. 2025. [Ontology-guided reverse thinking makes large language models stronger on knowledge graph question answering](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15269–15284, Vienna, Austria. Association for Computational Linguistics.
- Shihao Liu, Xiaofei Zhou, Bo Wang, and Geyuan Zhang. 2026. [Temporal evidence chain for temporal knowledge graph question answering with large language models](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 29515–29529, San Diego, California, United States. Association for Computational Linguistics.
- Vanessa Lopez, Michele Pasin, and Enrico Motta. 2005. [AquaLog: An ontology-portable question answering](#)

- system for the semantic web. In *The Semantic Web: Research and Applications*, volume 3532 of *Lecture Notes in Computer Science*, pages 546–562, Berlin, Heidelberg, Springer.
- Sviatoslav Lushnei, Dmytro Shumskyi, Severyn Shykula, Ernesto Jiménez-Ruiz, and Artur d’Avila Garcez. 2026. [Large language models as oracles for ontology alignment](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2435–2449, Rabat, Morocco. Association for Computational Linguistics.
- Yubo Ma, Zehao Wang, Mukai Li, Yixin Cao, Meiqi Chen, Xinze Li, Wenqi Sun, Kunquan Deng, Kun Wang, Aixin Sun, and Jing Shao. 2022. [MMEKG: Multi-modal event knowledge graph towards universal representation across modalities](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 231–239, Dublin, Ireland. Association for Computational Linguistics.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khoshabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Marshall McLuhan. 1964. *Understanding Media: The Extensions of Man*. McGraw-Hill, New York.
- Natasha Noy, Yuqing Gao, Anshu Jain, Anant Narayanan, Alan Patterson, and Jamie Taylor. 2019. [Industry-scale knowledge graphs: Lessons and challenges](#). *Communications of the ACM*, 62(8):36–43.
- OpenAI. 2026. GPT-5.6: Frontier intelligence that scales with your ambition. Model release report and system card.
- Michiel Overeem, Marten Spoor, Slinger Jansen, and Sjaak Brinkkemper. 2021. [An empirical characterization of event sourced systems and their schema evolution: Lessons from industry](#). *Journal of Systems and Software*, 178:110970.
- Palantir Technologies. 2026. [Action types: Overview](#). Palantir Foundry Documentation. Accessed 2026-07-26.
- Palantir Technologies Inc. 2022. [Ontology: Finding meaning in data](#). Palantir Blog. Accessed: 2026-07-24.
- Palantir Technologies Inc. 2026a. [Action types: Action log](#). Palantir Foundry Documentation. Accessed: 2026-07-24.
- Palantir Technologies Inc. 2026b. [Global branching: Overview](#). Palantir Foundry Documentation. Accessed: 2026-07-24.
- Palantir Technologies Inc. 2026c. [Ontology building: Overview](#). Palantir Foundry Documentation. Accessed: 2026-07-24.
- Kartik Sharma, Peeyush Kumar, and Yunqing Li. 2025. [OG-RAG: Ontology-grounded retrieval-augmented generation for large language models](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 32962–32981, Suzhou, China. Association for Computational Linguistics.
- Barry Smith, Michael Ashburner, Cornelius Rosse, Jonathan Bard, William Bug, Werner Ceusters, Louis J. Goldberg, Karen Eilbeck, Amelia Ireland, Christopher J. Mungall, Neocles Leontis, Philippe Rocca-Serra, Alan Ruttenberg, Susanna-Assunta Sansone, Richard H. Scheuermann, Nigam Shah, Patricia L. Whetzel, and Suzanna Lewis. 2007. [The OBO foundry: Coordinated evolution of ontologies to support biomedical data integration](#). *Nature Biotechnology*, 25(11):1251–1255.
- Yuan Sui, Yufei He, Nian Liu, Xiaoxin He, Kun Wang, and Bryan Hooi. 2025. [FiDeLiS: Faithful reasoning in large language models for knowledge graph question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8315–8330, Vienna, Austria. Association for Computational Linguistics.
- Theodore R. Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. 2023. [Cognitive architectures for language agents](#). *arXiv preprint arXiv:2309.02427*.
- Qingyun Sun, Jiaqi Yuan, Shan He, Xiao Guan, Haonan Yuan, Xingcheng Fu, Jianxin Li, and Philip S. Yu. 2025. [DyG-RAG: Dynamic graph retrieval-augmented generation with event-centric reasoning](#). *Preprint*, arXiv:2507.13396.
- Endel Tulving. 1972. Episodic and semantic memory. In Endel Tulving and Wayne Donaldson, editors, *Organization of Memory*, pages 381–403. Academic Press, New York.
- Zeno van Cauter and Nikolay Yakovets. 2024. [Ontology-guided knowledge graph construction from maintenance short texts](#). In *Proceedings of the 1st Workshop on Knowledge Graphs and Large Language Models (KaLLM 2024)*, pages 75–84, Bangkok, Thailand. Association for Computational Linguistics.
- Xiaoman Wang, Yaoze Zhang, Wenzhuo Fan, Hongwei Zhang, Ding Wang, Guohang Yan, Song Mao, Botian Shi, Yunshi Lan, and Pinlong Cai. 2026. [IA-RAG: Interval-algebra-driven temporal reasoning for dynamic knowledge retrieval](#). *Preprint*, arXiv:2606.06044.
- Shira Wein and Juri Opitz. 2024. [A survey of AMR applications](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6856–6875, Miami, Florida, USA. Association for Computational Linguistics.

Shira Wein and Nathan Schneider. 2024. [Assessing the cross-linguistic utility of Abstract Meaning Representation](#). *Computational Linguistics*, 50(2):419–473.

Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D. Manning, Percy Liang, and Jure Leskovec. 2022. [Deep bidirectional language-knowledge graph pretraining](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 37309–37323. Curran Associates, Inc.

Bowen Zhang and Harold Soh. 2024. [Extract, define, canonicalize: An LLM-based framework for knowledge graph construction](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9820–9836, Miami, Florida, USA. Association for Computational Linguistics.

Ze Yu Zhang, Zitao Li, Yaliang Li, Bolin Ding, and Bryan Kian Hsiang Low. 2026. [Respecting temporal-causal consistency: Entity-event knowledge graph for retrieval-augmented generation](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2017–2054, Rabat, Morocco. Association for Computational Linguistics.

Lingfeng Zhong, Jia Wu, Qian Li, Hao Peng, and Xindong Wu. 2024. A comprehensive survey on automatic knowledge graph construction. *ACM Computing Surveys*, 56(4).

Appendices

A Worked Example: Complementary Semantic Coverage Across Representations

This appendix illustrates the rationale behind Table 1 why KSP coordinates multiple knowledge representations rather than selecting a single universal representation. We use one sentence containing an entity reference, a nested event, temporal displacement, negation, attribution, and a previously unseen object:

On Tuesday, Maya claimed that the prototype drone SkyCart had not entered the restricted zone the previous night.

The comparison concerns each representation’s *native representational bias*, not its theoretical expressive limit. A sufficiently enriched knowledge graph can encode events, claims, and temporal qualifiers, just as an ontology can be populated with individuals and AMR graphs can be linked across documents. The relevant question is which representation makes each dimension explicit most directly, with the least additional modeling machinery.

A.1 Three Complementary Views

Knowledge-graph view. An enriched KG can represent the sentence by assigning stable identifiers to the entities, claim, and reported event:

```
:Maya          rdf:type :Person .
:SkyCart       rdf:type :PrototypeDrone .
:Zone17        rdf:type :RestrictedZone .
```

```
:Claim42       rdf:type :Claim ;
:speaker :Maya ;
:content :EntryEvent17 ;
:assertedAt :Tuesday .
```

```
:EntryEvent17 rdf:type :ZoneEntry ;
:agent :SkyCart ;
:destination :Zone17 ;
:polarity :Negative ;
:time :PreviousNight .
```

This representation is intentionally richer than a conventional subject–predicate–object extraction. A flat triple such as *Maya –claimed→ SkyCartEnteredZone17* would not adequately distinguish the claiming event from its propositional content. The enriched graph therefore reifies both the claim and the entry event. This preserves the distinction, but only because the KG has adopted an explicit event and assertion model.

AMR view. A simplified AMR graph directly represents the compositional structure of the sentence:

```
(c / claim-01
:ARG0 (p / person
:name (n / name :op1 "Maya"))
:time (d / date-entity
:weekday (t / tuesday))
:ARG1 (e / enter-01
:polarity -
:ARG0 (s / drone
:mod (p2 / prototype)
:name (n2 / name :op1 "SkyCart"))
:ARG1 (z / zone
:ARG1-of (r / restrict-01))
:time (n3 / night
:mod (p3 / previous))))
```

The outer predicate is the claiming event, with Maya as its speaker. Its :ARG1 is a separate entry event whose agent is SkyCart and whose destination is the restricted zone. Negation scopes over the entry event rather than the claim: Maya made a claim, but the content of that claim is that the

entry did not occur. The temporal expressions are likewise attached to the events they modify.

Ontology view. The ontology provides reusable concepts and constraints under which these instances can be interpreted:

PrototypeDrone	SubClassOf Drone
RestrictedZone	SubClassOf Zone
ZoneEntry	SubClassOf MotionEvent
Claim	SubClassOf CommunicationAct

hasSpeaker	Domain Claim
hasSpeaker	Range Person
hasContent	Domain Claim
hasContent	Range Proposition

hasAgent	Domain Event
hasAgent	Range PhysicalObject
hasDestination	Domain MotionEvent
hasDestination	Range Location

The ontology may additionally distinguish claims, observations, hypotheses, and verified facts, and define the evidence required to move information between these categories. For example, a system policy can specify that the content of a Claim must not automatically be promoted to an asserted world fact. Such a policy depends on the conceptual distinction between the communication act and the event it describes.

The three representations can be aligned through shared primitives. The AMR node for Maya links to the KG object :Maya, which instantiates the ontological class Person. The AMR enter-01 node links to :EntryEvent17, which instantiates ZoneEntry. The representations therefore describe the same objects and events while preserving different aspects of their meaning.

A.2 Factual Reference: Knowledge Graph

The sentence refers to three identifiable objects: Maya, SkyCart, and a particular restricted zone. The KG is the best fit for maintaining these referents as durable objects across sentences and documents. Once SkyCart has been assigned a stable identifier, subsequent mentions such as *the drone*, *the prototype*, or *it* can be linked to the same object. The graph can then accumulate properties and relations concerning SkyCart independently of the sentence in which it was first mentioned.

AMR represents the relevant participants, but its variables are initially local to the sentence. Additional entity linking or coreference resolution is

required to determine that the AMR node *s* denotes the same SkyCart represented elsewhere. The ontology can specify that SkyCart is a prototype drone, but an ontology primarily defines what kinds of objects may exist; it does not by itself provide durable memory for this particular drone. The KG is therefore strongest for preserving instance identity and factual reference over time.

A.3 Compositional Event Meaning: AMR

The sentence contains two nested events: Maya’s act of claiming and SkyCart’s alleged non-entry. Their participants, polarity, and temporal modifiers must remain correctly scoped. AMR expresses this composition directly. The entry event is the content of the claim, Maya is the participant in the claiming event, SkyCart is the participant in the entry event, and negation applies only to the latter.

A conventional KG extraction can easily flatten this structure. Recording only that Maya “claimed” an association between SkyCart and the zone obscures the internal event roles. Recording that SkyCart did not enter the zone without retaining the surrounding claim incorrectly treats the reported content as an established fact. An enriched KG can avoid these errors by reifying propositions and events, but doing so requires reconstructing much of the predicate–argument and scope structure that AMR provides directly.

The ontology supplies definitions for claims and entry events, but it does not determine how the predicates and arguments are composed in this particular utterance. AMR is therefore the best fit for preserving the sentence-local structure of who did what, to what, under which modifiers, and within whose scope.

A.4 Conceptual Abstraction: Ontology

The sentence presupposes conceptual distinctions that are not fully stated in its surface form. SkyCart is a drone and a physical object; a restricted zone is a kind of location; entering is a motion event; and claiming is a communication act rather than direct evidence that its content is true. These distinctions enable generalization beyond the individual sentence.

The ontology can encode inheritance, property restrictions, domain and range constraints, and equivalences among domain vocabularies. A query for all physical objects alleged to have entered controlled locations can therefore retrieve SkyCart even when the source text uses the more specific

terms *prototype drone* and *restricted zone*. The ontology can also detect incompatible assignments, such as representing a communication act as the physical agent of a motion event.

The KG may attach type labels to its objects, but those labels provide limited conceptual structure unless they are governed by an ontology or comparable schema. AMR identifies sentence-level concepts such as *drone*, *zone*, and *claim-01*, but it does not by itself define their domain-wide inheritance relations or constraints. The ontology is therefore the best fit for abstraction and conceptual reasoning.

A.5 Displacement and Temporality: AMR Event Ledger

The sentence describes an event displaced from the time of utterance. The claiming occurred on Tuesday, whereas the alleged non-entry occurred the previous night. AMR preserves both temporal expressions and attaches each to the event it modifies. This prevents *the previous night* from being incorrectly interpreted as the time of the claim.

AMR alone, however, does not necessarily resolve *Tuesday* or *the previous night* to calendar intervals. KSP therefore stores the AMR-derived events in an event ledger and resolves their relative times against document time, discourse context, and other events. Once Tuesday is anchored to a date, the previous night can be mapped to the corresponding interval. The ledger can then order the claim relative to the reported event and relate both to later observations.

A KG can store timestamps on event nodes, but temporal structure is not guaranteed by an ordinary entity–relation graph, and a current-state KG may discard the path by which its state was reached. An ontology can define intervals, temporal relations, and event categories, but it does not itself provide the sequence of observed events. The AMR-derived ledger is therefore best suited to capturing event-level temporal cues and organizing them into a history, while the KG materializes durable state and the ontology supplies the temporal vocabulary.

A.6 Modality, Negation, and Epistemic Status: AMR and Ontology

The sentence does not assert that SkyCart stayed outside the zone. It asserts that Maya made a claim whose content was that SkyCart did not enter. The distinction is essential: Maya may be correct, mis-

taken, deceptive, or reporting incomplete information.

AMR preserves the epistemic boundary structurally. The entry event remains embedded under *claim-01*, and its negative polarity remains attached to that embedded event. The ontology then classifies the outer event as a *Claim* and can associate it with a source, confidence, evidence status, or verification requirement. Together, these representations prevent the reported proposition from being silently promoted into the system’s asserted world state.

A KG can store the same distinction, as the enriched example does, but only if claims and propositions are explicitly reified. A plain triple such as *SkyCart -didNotEnter-> Zone17* would erase the fact that this information came from Maya’s statement. AMR is therefore strongest for extracting the relevant scope and polarity, while the ontology is strongest for determining how claims should be interpreted and governed.

A.7 Reflexiveness and Provenance: Ontology

The sentence is partly about another representation of the world: Maya’s statement. Answering questions such as “Who reported the non-entry?”, “What evidence supports it?”, or “Has the claim been independently verified?” requires the system to reason about statements and their sources as first-class objects.

AMR identifies the reporting event and its speaker, but it does not establish a general provenance model. The KG can attach source metadata or reify the statement, but the interpretation of that metadata depends on an explicit schema. The ontology can distinguish *Claim*, *Observation*, *Evidence*, *Source*, and *VerifiedAssertion*, and define the relations among them. It can also specify that a claim made by an agent is not equivalent to an observation produced by a sensor.

The ontology is therefore the best fit for provenance as a domain concept. The KG instantiates the provenance record, while AMR extracts the source relationship from the sentence.

A.8 Productivity: Extensible Capture and Controlled Schema Evolution

SkyCart may be a newly introduced prototype absent from both the KG and ontology. The KG can immediately add SkyCart as a new entity, while AMR captures its role in the sentence-level event,

including its name and *prototype* modifier, without requiring a predefined domain category.

This flexibility does not ensure shared semantics. The KG may lack an established type or relation for the new object, and AMR remains bounded by its general frames and relation inventory. The ontology can align SkyCart with an existing concept or introduce a class such as *PrototypeDrone*, but doing so requires validation and schema evolution.

KSP therefore treats productivity as a coordinated workflow: the KG admits new entities and facts, AMR preserves their event-level meaning, and the ontology determines which recurring distinctions should enter the shared conceptual model. Unresolved meanings can remain source-specific and provenance-linked until an appropriate alignment is established.

A.9 Combined Interpretation

No individual representation provides the complete interpretation. The KG establishes that Maya, SkyCart, and the restricted zone are durable objects. AMR reconstructs the two events, their participant roles, temporal modifiers, embedding, and negation. The ontology explains that a claim is a communication act whose content should not be treated as independently verified evidence.

KSP coordinates these contributions through aligned object and event identifiers. The sentence can therefore be stored without either flattening its compositional meaning or sacrificing durable retrieval. A later sensor observation that SkyCart did enter the zone would not overwrite Maya's earlier statement. Instead, the event ledger would retain both records, the ontology would distinguish the claim from the observation, and the KG would connect them to the same drone, zone, and time interval. The resulting knowledge state supports factual retrieval, conceptual reasoning, temporal reconstruction, and provenance-sensitive inference over a shared representation of the underlying objects and events.