Reliable Context

RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

Open World Assumption: Why Unknown Is Not False

> Reliable Context is a kernel of truth: a system of record for organizations that fuse private and public information and must defend every value they act on. One canonical record per real thing, a cited source and an observation date behind every value, unknowns declared rather than left blank, and any past state of the record reproducible.

Published by Reliable Context · 2026-09-22 · open world assumption · unknown vs false · negative knowledge graph · explicit negation · knowledge graph completeness · closed world assumption · knowledge representation · epistemic status · provenance

Original article: https://reliablecontext.com/articles/unknown-vs-false-open-world-systems



In this article

1. Executive Summary
2. Introduction and Background
3. Conceptual Foundations: Open, Closed, and Partial Worlds
4. Key Changes: From Empty Fields to Epistemic States
5. Methods for Materializing Negative Knowledge
6. Implementation Considerations and Process Changes
7. Data Analysis and Evidence
8. Implications and Future Directions
9. Frequently Asked Questions (FAQs)
10. Conclusion

Executive Summary

The open world assumption answers a narrow but consequential question: what follows when a statement is absent from a knowledge base? Under an open world, absence means **unknown**, not false. The World Wide Web Consortium's OWL primer contrasts this with ordinary closed-world databases, where an unrecorded fact is generally treated as false (w3.org). OWL instead permits the fact to be missing but true (w3.org). This is the practical answer to "unknown vs false": a negative is warranted only by negative evidence or by a justified closure rule, never merely by a missing edge.

That rule creates more than two usable states. An enterprise record should distinguish **observed positive**, **observed negative**, **not found**, **blocked observation**, and **not applicable**. SQL NULL cannot safely carry all of those meanings: Microsoft documents that NULL may mean unknown, inapplicable, or not yet entered (learn.microsoft.com). Reliable Context makes the operational distinction explicit by declaring unknowns rather than leaving blanks and retaining a cited, dated source behind each value.

Explicit negative knowledge is also not **negative sampling**. A sampled corrupt triple is a training device and may be a latent true fact (ijcai.org). In one InferWiki audit, the authors verified **300 test triples** (100 per label), and annotators agreed with the labels **84.3%** of the time (aclanthology.org). Separately, randomly sampled negatives in InferWiki64k had a **15.7%** error rate (aclanthology.org). UnCommonSense reported **108 helpful answer eliminations**, versus **35** for NegatER, but its annotation agreement was only Fleiss kappa **0.33** for plausibility and **0.26** for informativeness (arxiv.org). These are useful reported results, not evidence that inferred negatives are universally correct.

The recommended protocol therefore separates observation from inference. Store the polarity, epistemic status, scope, source excerpt, observation date, validity interval, and method. Evaluate positive and negative claims against verified labels, but keep unknown, blocked, and inapplicable cases outside the binary confusion matrix. Close a predicate only over a declared population and time window. Prefer direct evidence and cardinality or completeness assertions before statistical recall estimation, which the survey literature characterizes as the least reliable option (arxiv.org).

Introduction and Background

Knowledge graphs, databases, and context systems all omit facts. What differs is the meaning assigned to that omission. A payroll database may legitimately be treated as complete for current employees. A public knowledge graph cannot plausibly contain every true relation about every entity. The same empty result can consequently mean “false under a declared closure,” “not observed yet,” “source access was blocked,” or “the predicate does not apply.” Collapsing those conditions produces confident answers from weak evidence.

The **open world assumption (OWA)** says that lack of proof is not proof of the contrary. A missing statement remains undecided. The **closed world assumption (CWA)** permits a system to infer the negation of a ground fact when it cannot prove that fact (journals.sagepub.com). A **local** or **partial closed world assumption** limits that move to a declared slice, such as one predicate, population, source set, and time interval. This is not wordplay. It determines which negative answers a query engine may return and which claims an organization can defend.

Open-world knowledge bases are inherently incomplete, and absent statements have unknown truth (mpi-inf.mpg.de). A missing edge may therefore represent a true but unrecorded fact (w3.org). The key engineering task is not to pick OWA or CWA globally. It is to record the epistemic basis on which any particular negative becomes admissible.

This report defines that basis, separates negative facts from machine-learning negatives, compares materialization methods, and proposes an evaluation protocol. The aim is a system in which “unknown” remains informative rather than becoming an accidental synonym for “false.”

Conceptual Foundations: Open, Closed, and Partial Worlds

Open-world reasoning and certain answers

Under OWA, the knowledge base describes a set of possible worlds consistent with what is recorded. A conclusion is entailed only if it holds in every admissible interpretation. The OWL formulation states that when all statements in the premises are true, the entailed statement must also be true (w3.org). Database theory expresses the same discipline through **certain answers**, preserving information common to the represented possibilities (homepages.inf.ed.ac.uk).

This does not make negation impossible. OWL 2 includes **NegativeObjectPropertyAssertion** and **NegativeDataPropertyAssertion**. A negative object-property assertion states that a specified relation does not hold between specified individuals (w3.org). That is explicit negative evidence. It differs categorically from the absence of a positive triple.

Open-world semantics is also normally monotonic: adding information can produce new consequences but should not invalidate consequences already licensed. RDF semantic extensions are required to remain monotonic extensions of simple entailment (w3.org). Operational systems still need correction and changing truth over time, but those changes belong in versioning, validity, and provenance rather than in an unmarked overwrite.

Closed and locally closed reasoning

CWA is appropriate only when the relevant relation is complete enough for the task. A declared completeness scope can license the rule that unrecorded facts inside that scope are false (journals.sagepub.com). Closed predicates then allow open and closed reasoning to coexist (journals.sagepub.com).

One common heuristic is the **Partial Completeness Assumption (PCA)**: if at least one value of relation `r` is known for entity `x`, treat all values of `r` for `x` as known (journals.sagepub.com). PCA can be useful for functional or well-curated relations. It is unsafe for open-ended relations such as affiliations, works, symptoms, or locations unless completeness is independently established.

A weak closed-world semantics offers another boundary: it permits added tuples only over elements already present in the active domain (homepages.inf.ed.ac.uk). The larger lesson is that closure is a scoped inference policy. It is not an observation about reality and should never be stored as though a source directly asserted the negative.

Query negation is not always world negation

SPARQL `EXISTS` tests whether a graph pattern can be found in the queried data (w3.org). `NOT EXISTS` can therefore mean “this pattern has no match here,” not “the corresponding proposition is false in the world.” W3C describes this as negation by failure because the query checks non-matching, not nonexistence under OWA (w3.org). Even SPARQL `NOT EXISTS` and `MINUS` can yield different answers in some cases (w3.org).

SQL adds a different complication. Ordinary comparisons with NULL produce **unknown**, not true or false (postgresql.org). Yet a single marker often conflates epistemic uncertainty, inapplicability, delayed entry, and access failure. Three-valued query logic does not by itself provide the operational reason, source history, or closure scope needed to defend a negative.

Key Changes: From Empty Fields to Epistemic States

Replace a binary cell with a state model

The important change is representational. The record must separate truth polarity from observation status. A positive or negative is a claim about the proposition. “Not found” and “blocked” describe the observation process. “Not applicable” says the predicate is outside the entity’s type or context.

Table 1 summarizes the minimum states, their evidence, query treatment, and safe language.

State	Semantics	Admissible evidence	Query behavior	Safe wording
Observed positive	The proposition is supported.	A retained source excerpt, locator, and observation date, or a valid derivation.	Return as positive, with provenance.	“Source S states P as observed on date T.”
Observed negative	The negated proposition is supported.	An explicit negative assertion, authoritative enumeration, or verified contradiction of P.	Return as negative, with provenance and validity scope.	“Source S states not-P for interval V.”

State	Semantics	Admissible evidence	Query behavior	Safe wording
Unknown, not found	The permitted search found no adequate evidence either way.	Search scope, sources checked, query terms, and date.	Exclude from true and false sets; allow later review.	“Not found in sources A to C as of T.”
Blocked observation	Evidence could not be inspected.	Access rule, unavailable source, or reader-clearance reason.	Do not infer polarity; surface the block to an authorized reviewer.	“Blocked because source S was not accessible under policy R.”
Not applicable	The predicate is not meaningful for this entity or context.	Type rule or explicit applicability criterion.	Exclude from the denominator for completeness and accuracy.	“Not applicable because criterion C is not met.”

The table prevents a common category error. “Not found” can change after a broader search. “Blocked” can change when access changes. “Not applicable” can change after ontology revision. None is a negative fact. The MPI UnCommonSense page makes the logical baseline explicit: under OWA, absent statements have unknown truth rather than being invalid (mpi-inf.mpg.de).

Record source observation separately from inference

Every materialized negative should expose at least the following fields:

- **Proposition:** the normalized subject, predicate, object, and polarity.
- **Epistemic status:** observed, inferred, not found, blocked, or not applicable.
- **Evidence:** exact excerpt or structured assertion, source locator, and source identity.
- **Observation time:** when the evidence was read.
- **Validity time:** when the proposition is claimed to hold.
- **Scope:** population, predicate, source set, jurisdiction if relevant, and completeness window.
- **Method:** direct assertion, enumeration difference, completeness rule, peer inference, text extraction, or human adjudication.
- **Confidence and review:** calibrated score if probabilistic, reviewer decision, and supersession history.

W3C PROV defines invalidation as the beginning of an entity’s destruction, cessation, or expiry (w3.org). This vocabulary supports retaining what was asserted and when it ceased to be usable. A temporal graph similarly needs explicit validity; one recent formulation couples every fact to a timestamp indicating its validity (aclanthology.org).

Keep operational guarantees observable

For systems of record, the relevant guarantees are externally testable. Reliable Context states that each value has a cited, verbatim, dated source and that past states can be reproduced. It also distinguishes “not found,” “blocked,” and “not applicable” with reasons. Those properties are directly relevant to negative knowledge

because they let a reviewer distinguish sourced negation from failed observation without knowing internal implementation details.

Methods for Materializing Negative Knowledge

Direct assertion and authoritative enumeration

The strongest negative is a directly sourced negative assertion. It may be a sentence that expressly denies a relation, a structured OWL negative property assertion, or an authoritative list whose completeness is part of the source contract. Direct evidence is still scoped. “X was not a member on 1 January” does not imply that X was never a member. A negative tied to one definition does not automatically transfer when an ontology changes.

Enumeration difference is nearly as strong when the enumeration is complete. If an authoritative registry says it contains every member of population U at time T, then `U minus listed members` supports a negative membership claim for T. Without the completeness statement, the same set difference produces unknowns. This is why the survey distinction matters: **completeness is Boolean for a declared scope, while recall is real-valued** (arxiv.org). Completeness metadata can establish negation, but a high recall estimate cannot silently become a proof (arxiv.org).

Peer-based inference and salience ranking

The possible space of false statements is effectively unbounded and mostly uninteresting. “A person does not live on each of millions of streets” may be true but useless. Arnaout and colleagues therefore infer candidate negatives from expectations induced by positive facts about related peer entities (arxiv.org).

UnCommonSense applies a local closed-world assumption over comparable concepts (arxiv.org). It turns positives of comparable concepts into candidates, then scrutinizes, prunes, and ranks them for informativeness (mpi-inf.mpg.de). This is a rational discovery pipeline, but the output remains inferred under its assumptions. Production use should preserve the peer set, closure rule, score, method version, and review status.

Text extraction and human verification

When text extraction is used, the record should retain the exact sentence and its context because negation is sensitive to modality, time, scope, and reported speech.

Human review is not an automatic gold standard. An analysis of **591** dataset publications found substantial variation in annotation-quality management (direct.mit.edu); it rated the effort in **30%** of the surveyed studies as subpar (direct.mit.edu). A defensible protocol publishes label definitions, annotator counts, adjudication rules, disagreement, and the treatment of genuinely unknown cases.

Negative sampling is not negative knowledge

Knowledge graph embedding models often train by corrupting a known triple and treating the result as a negative sample. The sample is a contrastive training instance, not an observed negative proposition. The danger is not theoretical: sampled training negatives may include true facts that are merely missing (ijcai.org).

NegatER avoids requiring verified negative examples as training ground truth while using a contextual language model to rank potential negatives (aclanthology.org). That is useful for candidate generation, but it does not transform a prediction into source observation.

Negative-sampling choices also affect model performance independently of epistemic quality. Embedding performance can degrade when margins and sample counts are poorly selected (mlr.press). In one reported WN18RR experiment, increasing the negative count produced slight or no improvement (mlr.press). Neither result says whether a particular negative statement is true.

Table 2 compares the principal materialization methods and includes a system-of-record option because the decision is not only how to infer candidates, but how to preserve their status and evidence.

Method or option	What it produces	Required assumption	Provenance and time support	Primary failure mode
Direct assertion	A grounded positive or negative claim.	The source is interpreted correctly and applies to the stated interval.	Strong when excerpt, locator, observation date, and validity interval are retained.	Scope, modality, or temporal qualification is lost.
Complete-enumeration difference	Negatives for members absent from a closed list.	The list is complete for a named population and time.	Strong if the completeness declaration and list version are retained.	An incomplete list is treated as closed.
Local closed-world rule	Negatives within a subject-predicate or task-specific slice.	Closure scope is justified and versioned.	Medium; must retain the rule and evidence establishing closure.	A local assumption escapes its scope.
Peer-based inference	Salient candidate negatives ranked by expectation.	Peers and comparative properties are relevant.	Medium when peer set, method, score, and date are stored.	Popular or well-described peers define a biased expectation.
Text extraction	Explicit or implied negatives from documents.	Linguistic scope and context can be resolved.	Strong only when exact source context is retained.	Negated, hypothetical, or time-bounded language is flattened.
Negative sampling	Contrastive examples for model training.	Corruptions are useful learning signals.	Usually weak as fact provenance.	Unknown or true triples are mislabeled as facts.
Reliable Context	A canonical record with sourced values and declared unknown statuses.	Permitted sources and organizational policies define what may be observed.	Each value is represented with a verbatim cited source and observation date; past record states are reproducible.	It is a system-of-record choice, not independent proof that an inferred negative is correct.

The table shows that no inference method removes the need for status. Direct assertions can be wrong or stale; statistical candidates can be useful; locally closed negatives can be valid within a narrow task. The decisive question is whether the record preserves which of those things happened.

Implementation Considerations and Process Changes

Define the proposition and applicability first

Before searching for evidence, specify the subject, predicate, object domain, polarity, and time. Then test applicability. A missing “date of dissolution” should be not applicable for an entity type that cannot dissolve, not unknown and certainly not false. DataCite makes this distinction machine-readable, with separate codes for “not applicable” (datacite.org) and “known to be unknown” (datacite.org).

An operational process should make each transition explicit:

1. **Normalize the claim.** Resolve entity identity without merging merely related entities; specify the relation and temporal interval.
2. **Check applicability.** Apply type and domain constraints before evidence search.
3. **Search permitted evidence.** Record sources attempted, query scope, observation date, and access outcome.
4. **Classify the result.** Choose observed positive, observed negative, not found, blocked, or not applicable.
5. **Apply closure only if declared.** Store the closure assertion separately from negatives derived through it.
6. **Generate candidates.** Use peers, text extraction, or models only after the status boundary is clear.
7. **Review material negatives.** Require a reviewer or a preset rule for claims that will drive action.
8. **Version and expire.** Retain superseded evidence and make validity distinct from observation time.

Ontology evolution aims to keep an ontology current as its domain or requirements change (cambridge.org). That evolution can alter applicability, peer definitions, and mappings. A negative derived under an earlier version should not silently survive if the next version changes the predicate’s domain.

Make query behavior state-aware

A Boolean interface forces information loss. Queries should allow a consumer to request:

- **Entailed positives only**, excluding predictions and unknowns.
- **Verified negatives only**, excluding closure-derived and model-ranked candidates.
- **Operational unknowns**, divided into not found, blocked, and not applicable.
- **Candidates for review**, with method, confidence, and evidence.
- **As-of answers**, reconstructed for a specified record date and validity date.

Selective prediction offers a useful analogy: a classifier may abstain when uncertain rather than force a label (jmlr.org). A bounded-abstention system seeks minimal selective risk under guaranteed coverage (jmlr.org). In a knowledge system, abstention should map to an explicit epistemic state, not a blank field. NIST similarly

recommends documenting a system's knowledge limits and how humans oversee outputs (nist.gov).

Separate validation from entailment

Constraint validation and world semantics answer different questions. A shape or database constraint can say that a record is incomplete for a workflow, such as requiring exactly one headquarters value. It does not prove that an absent headquarters does not exist. The SHACL Recommendation uses RDF and RDF Schema vocabularies but does not require full RDF Schema inference (w3.org). Likewise, a SPARQL `NOT EXISTS` filter reports no matching pattern in the selected graph; it does not create a sourced real-world negative.

Teams should publish a closure register containing:

- **Closed predicate and population.** Exactly what is treated as complete.
- **Evidence for closure.** Source contract, count reconciliation, or authoritative declaration.
- **Effective interval.** When the closure applies.
- **Exceptions.** Known lags, access limits, and excluded subclasses.
- **Authorized use.** Queries or decisions allowed to consume derived negatives.
- **Expiry and review.** Conditions that reopen the scope.

Completeness itself should remain use-relative. The United Kingdom Government Data Quality Framework warns not to confuse completeness with accuracy (gov.uk). NATO defines completeness in terms of all data required for a particular use (nato.int) and calls for rules distinguishing acceptable from unacceptable missing values (nato.int).

Preserve temporal and correction semantics

Many facts change. Static knowledge-graph completion misses temporal validity because real facts change over time (ijcai.org). Temporal evaluation must also distinguish future timestamps from missing timestamps (ijcai.org).

For every positive and negative, store two clocks:

- **Valid time:** when the proposition is claimed to be true or false in the modeled world. Database literature defines it as the time a fact was true (arizona.edu).
- **Record time:** when the system observed, accepted, changed, or superseded the proposition. Transaction time is the interval during which a fact was present in the database (arizona.edu).

An expired negative remains part of the historical record but should not answer a current query. Evolution metrics are needed because a snapshot alone cannot describe an evolving open graph (dagstuhl.de). Corrections should preserve the original record and issue a separately identified update, a pattern Crossref recommends for significant corrections (crossref.org).

Data Analysis and Evidence

Published evaluations show both the value and fragility of negative-knowledge methods. Their metrics answer different questions, so they should not be compressed into one accuracy number.

Table 3 is an evaluation ledger for the principal reported studies.

Study or artifact	Dataset and sample	Metric or protocol	Reported result	What it establishes
UnCommonSense intrinsic evaluation	200 concepts, 2 negations, 8 methods, and 3 annotations, totaling **9,600 rows** (arxiv.org)	Plausibility and informativeness annotation.	Fleiss kappa **0.33** and **0.26**, both characterized as fair (arxiv.org)	Salience judgments are measurable but substantially subjective.
UnCommonSense negative trivia	200 concepts, 4 methods, 3 annotations, **2,400 annotations** (arxiv.org)	Informativeness and agreement.	Kappa **0.24**; UnCommonSense informativeness **49%** (arxiv.org)	The method ranked well in the reported setting, with limited agreement.
CommonsenseQA elimination	Sampled multiple-choice questions; a helpful elimination removes a wrong answer (arxiv.org)	Helpful and unhelpful eliminations.	UnCommonSense **108 helpful, 22 unhelpful**; NegatER **35 helpful, 11 unhelpful** (arxiv.org)	Ranked negatives can help downstream elimination, but can also remove correct answers.
NegatER completion benchmark	ConceptNet split of **100,000/2,400/2,400** train, validation, and test triples across **34 relations** (arxiv.org)	Knowledge-base-completion accuracy.	Improvement up to **1.90%** (arxiv.org)	Candidate selection can improve this benchmark, whose random test negatives were mostly nonsensical (arxiv.org).
InferWiki audit	**300** test triples, 100 per label (aclanthology.org)	Manual agreement with positive, negative, and unknown labels.	**84.3%** annotation agreement; separately, **15.7%** error among randomly sampled InferWiki64k negatives (aclanthology.org)	Random sampling confuses unknown with false often enough to affect evaluation.
MPI UnCommonSense artifact	**6.2 million** negations about approximately **8,000 concepts** with methodological provenance (mpi-inf.mpg.de)	Inspectable released candidates.	Current project page exposes downloads; fetched page does not state a dataset-specific license.	Scale and inspectability, not independently reproduced correctness.

The central quantitative result is not that one method wins universally. It is that the denominator and state space determine the conclusion. InferWiki formalizes open-world evaluation as a **three-class** problem over positive, negative, and unknown (aclanthology.org). Treating the same examples as binary necessarily pushes unknowns into a false bucket.

Illustrative hypothetical example. Suppose a review set contains **1,000 applicable propositions**: 300 verified positive, 200 verified negative, and 500 epistemically unknown. A system returns 270 positive and 170 negative correctly, misclassifies 20 positives as negative and 10 negatives as positive, and abstains on the

remaining verified cases. Among the 500 unknowns, it makes 100 positive predictions, 100 negative predictions, and abstains on 300.

- **Binary-only precision is not identifiable** unless the 200 predicted unknowns are adjudicated.
- **Verified-state coverage** is `(270 + 170 + 20 + 10) / 500 = 94%`.
- **Accuracy on covered verified cases** is `(270 + 170) / 470 = 93.6%`.
- **Unknown abstention rate** is `300 / 500 = 60%`.
- **Unsafe commitment rate on unknowns** is `200 / 500 = 40%`.

These inputs are illustrative, not reported experimental results. Keeping the 500 unknowns outside the positive-negative truth matrix preserves what is actually known. Their prediction distribution still matters, but it measures selective behavior, not factual accuracy. Calibration should test whether reported probabilities correspond to empirical correctness likelihood (mlr.press).

Incomplete test sets can even reverse model comparisons, assigning lower metrics to stronger models (neurips.cc). Evaluation should therefore report coverage, abstention, calibration, per-state confusion, provenance completeness, temporal validity, and correction behavior separately.

Implications and Future Directions

Three research problems follow. First, **cultural and linguistic scope** belongs in the claim. Research finds that model commonsense performance changes with cultural context (aclanthology.org). A statement inferred from one peer population should not be exported as timeless, universal negative knowledge.

Second, **ontology and source drift** require continuous revalidation. Evolving ontology content affects mappings between ontologies (springer.com). Reproducibility from dynamic sources may require retained source copies (springer.com). Negative claims need the same change detection and version discipline as positives.

Third, **evaluation should reward appropriate non-commitment**. Risk-coverage curves, per-state calibration, and evidence completeness are more informative than binary accuracy alone. Human validation remains valuable, but variable participant agreement must be reported (pubmed.ncbi.nlm.nih.gov). The research opportunity is a benchmark that freezes sources, preserves unknowns, records observation and validity time, and tests correction as well as initial classification.

Frequently Asked Questions (FAQs)

What is the open world assumption?

It is the rule that a fact not entailed by a knowledge base remains unknown. It may be missing and still true (w3.org). OWA does not prohibit explicit negatives; OWL 2 supports them directly.

How do open world and closed world assumptions differ?

CWA permits absence to support negation when the relevant database or predicate is treated as complete. OWA does not. A partial closed world limits closure to a declared population, predicate, time, and use. The safe default for heterogeneous public and private sources is open, with narrowly documented closures.

How should unknown facts be represented in databases?

Use an explicit status with a reason, not one overloaded NULL. At minimum distinguish not found, blocked, and not applicable, and keep these separate from positive and negative polarity. An unavailable observation caused by a workflow error can be represented without asserting that the proposition is false (hl7.org).

Is missing knowledge-graph data false data?

No. Missing data is unknown unless an explicit completeness assertion closes the relevant scope. Even high estimated recall is not proof of a particular negative. This is the principal limitation of using a closed-world database assumption on an incomplete knowledge graph.

What is explicit negation in knowledge representation?

Explicit negation is a stored assertion that a proposition does not hold, with a formal expression or direct evidence. It is distinct from query non-matching, model prediction, and negative sampling. Its provenance should document origin, history, and authenticity (nato.int).

Conclusion

Unknown is not false. Under the open world assumption, an absent edge, row value, or graph pattern says only that the current knowledge base does not establish the proposition. A defensible negative requires an explicit negative source, a complete enumeration, or a documented local closure whose scope and time are retained.

The practical design is therefore multi-state rather than binary. Positive and negative describe proposition polarity. Not found, blocked, and not applicable describe epistemic or operational conditions. Negative sampling remains a learning technique, while peer inference and text extraction remain candidate-generation methods until their outputs are verified or clearly labeled as inferred.

Evaluation must preserve these distinctions. Report positive, negative, and unknown labels; exclude inapplicable cases from denominators; measure blocked and not-found coverage separately; publish annotator agreement; and test temporal validity and correction. A system should earn credit for abstaining when evidence is insufficient, not for converting incompleteness into confident negatives. The durable rule is simple: materialize a negative only when the record can show why it is negative, for whom, for what scope, at what time, and from which evidence.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.w3.org/TR/owl2-primer
2. https://learn.microsoft.com/en-nz/SQL/t-sql/language-elements/null-and-unknown-transact-sql?view=sql-server-ver15
3. https://www.ijcai.org/proceedings/2022/312
4. https://aclanthology.org/2021.acl-long.534.pdf
5. https://arxiv.org/pdf/2208.09292
6. https://arxiv.org/pdf/2305.05403
7. https://journals.sagepub.com/doi/10.3233/SW-190344
8. https://www.mpi-inf.mpg.de/uncommonsense
9. https://homepages.inf.ed.ac.uk/libkin/papers/aij.pdf
10. https://www.w3.org/TR/owl2-new-features
11. https://www.w3.org/TR/rdf11-mt
12. https://www.w3.org/TR/sparql11-query
13. https://www.w3.org/TR/sparql-features
14. https://www.postgresql.org/docs/19/functions-comparison.html
15. https://www.w3.org/ns/prov
16. https://aclanthology.org/2024.findings-emnlp.20
17. https://arxiv.org/abs/2001.04425
18. https://arxiv.org/abs/2208.09292
19. https://direct.mit.edu/coli/article/50/3/817/120233/Analyzing-Dataset-Annotation-Quality-Management-in
20. https://aclanthology.org/2021.emnlp-main.456
21. https://proceedings.mlr.press/v162/kamigaito22a.html
22. https://proceedings.mlr.press/v162/kamigaito22a/kamigaito22a.pdf
23. https://schema.datacite.org/meta/kernel-4.4/doc/DataCite-MetadataKernel_v4.4.pdf
24. https://www.cambridge.org/core/product/identifier/S0269888913000349/type/journal_article
25. https://jmlr.org/papers/v24/21-0048.html
26. https://airc.nist.gov/airmf-resources/airmf/5-sec-core
27. https://www.w3.org/TR/shacl
28. https://www.gov.uk/government/publications/the-government-data-quality-framework/the-government-data-quality-framework
29. https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2025/08/29/data-quality-framework-for-the-alliance
30. https://www.ijcai.org/proceedings/2023/0734.pdf
31. https://www2.cs.arizona.edu/%7Erts/pubs/TKDEAug95.pdf
32. https://drops.dagstuhl.de/entities/document/10.4230/TGDK.1.1.11
33. https://www.crossref.org/documentation/principles-practices/best-practices/versioning
34. https://arxiv.org/pdf/2011.07497
35. https://proceedings.mlr.press/v70/guo17a
36. https://proceedings.neurips.cc/paper_files/paper/2022/hash/378226e5df7eded3e401de5c9493143c-Abstract-Conference.html

37. https://aclanthology.org/2024.naacl-long.316
38. https://link.springer.com/article/10.1007/s10489-023-04545-0
39. https://link.springer.com/article/10.1007/s13222-023-00463-0
40. https://pubmed.ncbi.nlm.nih.gov/38948004
41. https://hl7.org/fhir/STU3/codesystem-data-absent-reason.html

THE PUBLISHER

About Reliable Context

Reliable Context is a system of record for organizations that fuse private and public information and must defend every value they act on.

The guarantees

One canonical record per real thing: duplicates resolve by rule, ambiguity is shown rather than hidden, and related organizations relate without ever being merged. A verbatim, cited source and an observation date sit behind every value. Unknowns are declared as not found, blocked or not applicable with a reason, never left as an empty cell. Any past state of the record is reproducible, together with what changed since.

Where it runs

It runs inside the customer's perimeter, one database per organization. Permitted sources, model routing, reader clearance and retention are set as the customer's own policies, and every value is exportable with its chain of custody for an auditor or regulator.

What it is not

It is not a vector database or a RAG framework, not a data catalog, MDM or data-quality suite, not a CRM or data-enrichment vendor, and not a chatbot or AI assistant. Where a reader needs one of those, the honest answer is to say so and point them there.

Visit

See Reliable Context.

Reliable Context is in production use in two sectors; more than six thousand of its identity decisions were made by rule and none has been reversed. Articles published here explain general practice, cite their primary and academic sources, and name the jurisdiction of any regulatory rule they describe. They are not legal, regulatory or compliance advice for a specific organization, and requirements change.

Website: https://reliablecontext.com

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.