Reliable Context

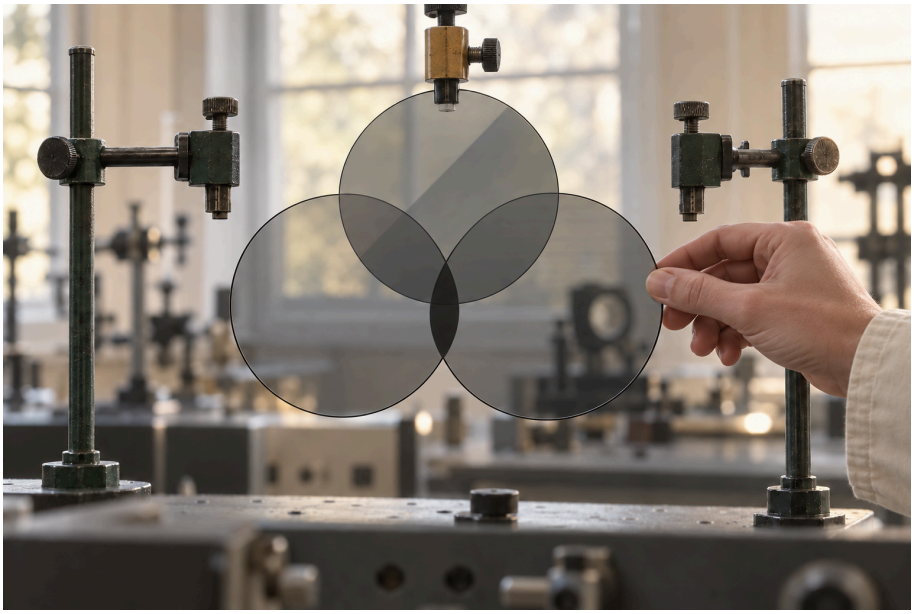
RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

Ontology Crosswalk Acceptance Tests for Mapping Quality

> 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 · ontology crosswalk acceptance test · ontology mapping validation · skos · semantic crosswalk · exactmatch · ontology alignment · sssom · mapping quality · knowledge governance

Original article: https://reliablecontext.com/articles/ontology-crosswalk-acceptance-test



In this article

1. Executive Summary
2. Introduction and Background
3. Key Changes
4. Predicate Selection and Formal Consequences
5. Acceptance Test Procedure
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

An **ontology crosswalk acceptance test** should treat every row as a versioned claim with operational consequences, not as a similarity score. The reviewer begins with the consumer action, such as retrieval, normalization, transformation, aggregation, or inference, and asks whether the asserted predicate preserves the distinctions that action requires. In Simple Knowledge Organization System (SKOS), mapping links conventionally connect concepts in different schemes (www.w3.org). `skos:exactMatch` is symmetric and transitive, while `skos:closeMatch` is symmetric but deliberately not transitive (www.w3.org). An acceptance test must therefore inspect what a predicate entails after closure, not merely whether two labels appear similar.

The minimum test packet pins both source releases and preserves the reviewed identifiers, labels, definitions, predicate, justification, authorship, review, observation date, status, and evidence. **SSSOM 1.0.0** requires subject, object, predicate, and mapping justification on each mapping (mapping-commons.github.io), while its version fields can hold a version Internationalized Resource Identifier (IRI) or version string (github.com). The current FAIR Mapping technical guidance recommends SSSOM as a sharing model, but the Research Data Alliance (RDA) working-group output page listed no formal outputs as of **September 22, 2026** (mapping-commons.github.io) (www.rd-alliance.org). RDA reserves “RDA Recommendation” for official, endorsed group results (www.rd-alliance.org). The online material should therefore be described as evolving technical guidance, not an endorsed RDA Recommendation.

The black-box procedure has **nine controls**: define the consumer decision; pin releases; validate identifiers and scheme membership; reject duplicates and predicate clashes; select the weakest adequate relation; materialize symmetry, inverses, subproperties, and transitivity; run positive, negative, and boundary queries; simulate version drift; and issue an approval, downgrade, rejection, or defer decision. “No acceptable mapping” is a successful reviewed outcome when every available predicate overstates the evidence. Confidence may prioritize review, but a number between 0 and 1 is not automatically a calibrated probability, even though SSSOM permits such a score (github.com).

Evidence supports this conservative design. Beisswanger and Hahn found **30 erroneous correspondences** that required removal from a reference alignment (pmc.ncbi.nlm.nih.gov). A competency-question resource contains **234 requirements and 131 formal queries**, showing how intended behavior can be tested as

executable expectations (pmc.ncbi.nlm.nih.gov). The decisive output is not a universal pass rate. It is a reproducible decision for each row, with its evidence, consequences, reviewer, date, and re-review trigger.

Introduction and Background

An ontology crosswalk connects semantic entities governed in separate vocabularies. A **concept scheme** is the vocabulary context in which a concept is identified and related. A **semantic entity** is the subject or object being mapped. A **correspondence** is the proposed relation between those entities. A **mapping set** is the governed collection of correspondences. A **mapping predicate** states the semantic commitment. A **justification** records why that commitment is warranted. A **source version** identifies the release whose labels, definitions, and structure were reviewed.

These distinctions matter because the row may be consumed far from the review interface. A search service may expand a query through an exact mapping. A transformation pipeline may replace one identifier with another. An aggregate may combine observations under a broader category. A reasoner may materialize symmetric, inverse, or transitive consequences. The Library of Congress Linked Data Service illustrates the identifier baseline: its lists and individual values have resolvable Uniform Resource Identifiers (URIs) (www.loc.gov). UK government guidance connects persistent identifiers to the practical ability to combine data from different sources (www.gov.uk).

The acceptance question is therefore specific: **can this predicate, between these two versioned entities, support these named consumer actions without producing an unacceptable positive, negative, or inferred result?** That is narrower than evaluating a mapping generator. It is also stronger than checking label similarity. Published work on user validation concludes that automatically produced alignments require manual validation when interoperability is the intended use (www.cambridge.org).

In search terms, the method is an **ontology mapping validation** protocol, a set of **semantic crosswalk quality criteria**, a **hierarchical ontology mapping** check, a **SKOS mapping relation validation** procedure, and an **ontology crosswalk review checklist**. The sections below answer each of those intents through one consequence-based method rather than separate glossaries.

For organizations that must defend mapping-derived values, the accepted row also needs a durable evidence record. **Reliable Context** documents one canonical record per real thing, a cited and dated source behind each value, declared unknowns, and reproducible past states. Those are relevant observable properties for a mapping decision register. They do not replace a terminology editor, reasoner, or test harness.

Key Changes

From a similarity output to a claim under test

The first change is epistemic. A score may explain why a candidate reached a reviewer, but it does not define the accepted meaning. The mapping predicate does. SSSOM distinguishes the creator that publishes a mapping in SSSOM form from the author who asserted it (academic.oup.com), and it separately represents reviewers (academic.oup.com). Acceptance should likewise separate proposal, assertion, review, and publication roles.

The review begins by stating the failure cost. Retrieval often tolerates controlled expansion that transformation cannot. Aggregation may accept a broader category if direction is preserved. Identity resolution requires evidence beyond a concept-mapping predicate. Logical inference has the strictest need to inspect closure because a locally plausible assertion can create remote consequences.

From one positive pair to a boundary fixture

A row is not accepted after one confirming example. The reviewer constructs a small fixture that makes prohibited outcomes visible:

- **Positive query:** an expected source item must find or transform to the target.
- **Negative query:** a confusable neighbor must remain excluded.
- **Boundary query:** a peripheral or polysemous item tests where interchangeability stops.
- **Directional query:** broad and narrow mappings must behave differently when traversed.
- **Aggregate query:** counts must not be duplicated through multiple paths.
- **Inference query:** every entailed pair must be inspected, not only asserted pairs.

Competency questions offer a precedent for executable expectations. A published resource tested formalized queries to verify that an ontology yielded expected answers (pmc.ncbi.nlm.nih.gov). The fixture here is intentionally black-box: it evaluates observable results, not how candidate mappings were generated.

From timeless rows to release-bound decisions

The accepted claim is bound to two releases. A label can remain stable while a definition, parent, status, or extension changes. W3C data best practice recommends a complete history explaining changes in every version (www.w3.org). The EU Publications Office provides a concrete stability pattern: during one label restructuring, the current concept URI remained identical (op.europa.eu). Identifier stability is useful, but it does not prove semantic stability.

Version review must classify each affected row as unchanged, still valid, downgraded, split, merged, retired, or unknown. The Unified Medical Language System (UMLS) distributes change files so consumers can determine whether changes affect their applications (www.ncbi.nlm.nih.gov). Its documentation also warns that splits are not explicitly represented in one concept-history file, requiring analysis of atom movements (www.nlm.nih.gov). The EU Publications Office separately cautions that downloaded authority-table files can become deprecated while still being used (op.europa.eu). A release number alone is therefore insufficient evidence of continued validity.

Predicate Selection and Formal Consequences

Predicate selection should use the weakest relation that safely supports the consumer action. It should not force a mapping where evidence supports none. SKOS `inScheme` supplies explicit scheme context because its range is `skos:ConceptScheme` (www.w3.org). This matters because SKOS mapping properties are intended for cross-scheme links, not ordinary within-scheme hierarchy.

Table 1 is the predicate-selection matrix. “Safe use” means safe only after the row passes the consumer-specific fixture.

Outcome	Allowed subject and object	Semantic commitment and inference	Potentially safe consumer use	Reject when
`skos:exactMatch`	Two SKOS concepts in different schemes	High-confidence broad interchangeability; symmetric, transitive, and a subproperty of `closeMatch` (www.w3.org)	Broad retrieval expansion after closure testing	Any boundary case differs, or transitive closure introduces an unsafe pair
`skos:closeMatch`	Two SKOS concepts in different schemes	Similar enough for some retrieval uses; symmetric and non-transitive (www.w3.org)	Explicitly scoped query expansion	A consumer treats it as substitution, identity, or transitive equivalence
`skos:broadMatch`	Specific source concept to broader target concept	Directed hierarchical mapping; object is more generic (www.w3.org)	Roll-up or recall-oriented navigation	Direction is reversed, or aggregation would double-count
`skos:narrowMatch`	Broad source concept to narrower target concept	Inverse direction of `broadMatch`; not a claim of equivalence	Drill-down candidate generation with explicit direction	The target is only associated, not genuinely narrower
`skos:relatedMatch`	Associatively related concepts in different schemes	Symmetric associative link, not transitive (www.w3.org)	Discovery and related-item navigation	A consumer uses it for replacement, roll-up, or logical identity
`owl:sameAs`	Individuals that are the same resource	Full identity between individuals (www.w3.org)	Identity reasoning only when identity has actually been established	The endpoints are merely similar concepts or classes
`owl:equivalentClass`	OWL classes with equal extensions	Same class extension (www.w3.org)	Class-level reasoning after equivalence proof	The endpoints are SKOS concepts or only retrieval-interchangeable
No acceptable mapping	Any proposed pair	No predicate meets the stated consequence threshold	Preserve uncertainty and prevent unsupported propagation	Never reject merely because a downstream system expects complete coverage

The matrix separates **concept mapping**, **individual identity**, and **class equivalence**. Substituting one for another changes the claim, not just the syntax. It also makes downgrade meaningful: an exact mapping can become close, hierarchical, related, or unmapped without erasing the reviewed evidence.

Two integrity rules deserve explicit automated tests. First, `exactMatch` is disjoint with `broadMatch` and `relatedMatch` (www.w3.org). Through inverse consequences, a clash with `narrowMatch` must also be detected. Second, the only SKOS mapping relation declared transitive is `exactMatch`. A system must not infer A closeMatch C merely because A closeMatch B and B closeMatch C.

Acceptance Test Procedure

The procedure below evaluates a proposed mapping set without inspecting its generation algorithm. It is suitable for a sample during procurement, for every high-impact row, or for every row when the set is small enough.

Establish the evidence packet

The reviewer first freezes what was observed. SSSOM's mandatory core includes the two identifiers, predicate, and justification (mapping-commons.github.io). The FAIR Mapping guidance recommends creator and reviewer identifiers, versions, confidence, and mapping-set metadata as mandatory operational fields (mapping-commons.github.io). Dublin Core similarly recommends a URI for identifying a creator (www.dublincore.org).

Table 2 defines a reviewable evidence row. It is a minimum acceptance artifact, not a claim that one serialization is universally required.

Field group	Required content	Acceptance question
Endpoints	Subject ID, target ID, scheme membership, resolvable URI	Do both identifiers resolve to the exact entities reviewed?
Observed meaning	Preferred labels, definitions, language, status, relevant parents and children	Is the decision grounded in meaning and structure rather than labels alone?
Versions	Subject release, target release, retrieval or observation date	Can another reviewer reconstruct the same comparison?
Claim	Predicate, direction, mapping-set ID and version	Is the semantic commitment explicit and correctly directed?
Rationale	Justification type, concise reasoning, evidence URL and locator	Does the evidence warrant the predicate rather than merely nominate the pair?
Responsibility	Creator, author where distinct, reviewer, decision owner	Is each role attributable through a persistent identifier where available?
Decision	Approved, downgraded, rejected, or deferred; reason code	Is uncertainty represented as a governed outcome rather than a blank?
Lifecycle	Observation date, next review or event trigger, prior state	Can a consumer determine what changed and when to retest?
Test evidence	Fixture version, assertions, inferred pairs, query results	Were actual downstream consequences exercised?

The table makes absence explicit. A missing version can be marked **not found**; an inaccessible source can be **blocked**; an inapplicable predicate field can be **not applicable**, each with a reason. Reliable Context documents this distinction as “not found, blocked, or not applicable,” rather than emitting an empty cell. That pattern is directly relevant to an evidence register because a blank cannot distinguish missing research from a negative finding.

Run mechanical and semantic checks

The acceptance sequence is:

1. **Resolve endpoints.** Dereference both identifiers and record the observed response.
2. **Confirm scheme membership.** Verify that each entity belongs to the claimed source or target scheme.
3. **Pin releases.** Record immutable version identifiers where available, otherwise the publisher's release label and observation date.
4. **Compare meaning.** Preserve labels and definitions, then inspect scope notes, exclusions, hierarchy, status, and language.
5. **Detect duplicates.** Reject duplicate rows, reversed duplicates where symmetry makes them redundant, and multiple active statuses for one claim.
6. **Detect predicate clashes.** Materialize subproperty, symmetry, inverse, and transitivity rules before checking disjointness.
7. **Test direction.** For hierarchical mappings, verify which endpoint is broader and which is narrower.
8. **Execute the fixture.** Run positive, negative, boundary, directional, aggregate, and inference queries.
9. **Issue the decision.** Approve, downgrade, reject, or defer, with owner, observation date, reason, and trigger.

Sun and colleagues published open N3 validation rules for SKOS mapping quality (arxiv.org). Their rules include patterns associated with SKOS vocabulary hijacking (arxiv.org). Automated rules are valuable for repeatable rejection conditions, but they do not answer whether the accepted row supports a particular consumer decision.

Inspect closure, not only assertions

For every asserted pair, materialize these consequence classes in an isolated graph:

- **Symmetry:** reverse `exactMatch`, `closeMatch`, and `relatedMatch` links.
- **Inverse:** derive `narrowMatch` from `broadMatch`, and the reverse.
- **Subproperty:** treat `exactMatch` as a `closeMatch` consequence.
- **Transitivity:** chain `exactMatch` only.
- **Disjointness:** flag exact, broad, narrow, or related combinations that cannot coexist.
- **Duplicate paths:** record whether several assertions yield the same entailed pair.

OAEI evaluation practice supports relation-aware and logical evaluation. Its Conference track includes reasoning checks based on conservativity-principle violations (oaei.ontologymatching.org). Its Beyond Equivalence track gives partial credit to superclass, subclass, and overlap relations rather than collapsing every valid relation to equivalence (oaei.ontologymatching.org). An acceptance test should be at least as explicit about relation type.

Implementation Considerations and Process Changes

Govern decisions, not just files

Four statuses make review operational:

- **Approved:** the asserted predicate and all tested consequences meet the named consumer threshold.
- **Downgraded:** the pair is useful, but only under a weaker predicate or narrower consumer scope.
- **Rejected:** evidence contradicts the proposed predicate, a mechanical rule fails, or the fixture exposes an unsafe result.
- **Deferred:** evidence is missing, blocked, conflicting, or awaiting a pinned release. The row is not active.
- **No acceptable mapping:** a reviewed subtype of rejection or approval-to-remain-unmapped, depending on local workflow.

The last outcome is essential. Coverage is not quality. A complete-looking crosswalk can be worse than a sparse one if consumers cannot distinguish supported mappings from forced ones.

Ownership should also be explicit. Getty subjects contributed vocabulary data to an established review procedure and retains final editorial control (www.getty.edu). This separates participation from decision authority. An internal crosswalk should likewise name a decision owner even when evidence and proposals come from several teams.

Retest on semantic events

A calendar review is useful, but event triggers are more precise:

- **Release published:** either endpoint scheme issues a new release.
- **Definition changed:** scope, inclusion, or exclusion text changes.
- **Hierarchy changed:** a parent, child, or grouping changes.
- **Status changed:** an endpoint becomes deprecated or retired.
- **Entity split or merge:** one concept becomes several, or several consolidate.
- **Predicate policy changed:** the organization changes the allowed use of a relation.
- **Consumer changed:** a mapping approved for retrieval is proposed for transformation or reasoning.
- **Failure observed:** a production query yields a prohibited positive, negative, or aggregate.

UMLS retains provenance while enabling cross-terminology mappings (www.nlm.nih.gov). It also maps retired identifiers to synonymous or related current identifiers (www.ncbi.nlm.nih.gov). Those are distinct outcomes and should not be flattened into unchanged equivalence.

Choose an operational record

The mapping file and the decision record may be the same artifact, or the decision evidence may live in a governed system linked to an exported mapping set. *Table 3* is the primary operational comparison. It compares evidence-record options, not mapping algorithms.

Option	Useful when	Acceptance evidence and limitations
Versioned SSSOM mapping set	Interchange and machine validation are primary	Strong structured representation for identifiers, predicates, justification, versions, confidence, and review metadata. **SSSOM 1.0.0** remains the stable specification (github.com); a **1.1.0 alpha 5** pre-release existed in May 2026 (github.com). A file still needs policy, test execution, and durable source evidence.
Vocabulary-management workflow	Editors already govern concepts, releases, and contributions	Can align mapping approval with editorial authority. Confirm that it preserves old states, external evidence, test fixtures, and exportable decisions rather than assuming those properties. BioPortal submission metadata, for example, includes alpha, beta, production, and retired lifecycle states (www.bioontology.org).
Reliable Context system of record	Mapping-derived values must be defended alongside fused private and public information	The vendor documents verbatim cited evidence, observation dates, declared unknowns, replayable past states, customer-set source and retention policies, and export with chain of custody. It is an evidence-record option, not a substitute for predicate semantics or a mapping test harness.

The appropriate choice depends on the downstream record obligation. A portable mapping set may be sufficient for public interchange. A vocabulary workflow may be best for editorial control. A system of record becomes relevant when the organization must reproduce which evidence and state supported a consequential mapping-derived value at a past date. These options can also be complementary.

Data Analysis and Evidence

Closure calculation on a synthetic chain

Consider three concepts, **A**, **B**, and **C**, in three independently governed schemes. The asserted set contains two rows: A `exactMatch` B and B `exactMatch` C. Because exact mapping is symmetric and transitive, the closure contains all six directed pairs among the three concepts, plus the corresponding `closeMatch` consequences. The two assertions therefore create **four newly implied directed exact pairs**: B to A, C to B, A to C, and C to A. The surprising pair is often A to C because nobody reviewed it directly.

The same chain with `closeMatch` produces only the two reversed pairs, not A to C or C to A, because close mapping is not transitive. This difference is the operational reason to test predicate consequences. W3C defines exact mapping for interchangeability across a wide range of information-retrieval applications (www.w3.org), while deliberately withholding transitivity from close mapping.

The closure test records:

- **Asserted exact pairs:** 2.
- **Directed exact pairs after closure:** 6.
- **New exact pairs requiring inspection:** 4.

- **Directly unreviewed endpoint pair:** A and C.
- **Close-match consequences:** 6, because exact mapping is a close-mapping subproperty.

These are arithmetic results for the synthetic fixture, not a universal defect rate. The acceptance criterion is that every newly implied pair passes the same negative and boundary expectations as a directly asserted row.

What published evaluation numbers do and do not show

Beisswanger and Hahn proposed **10 quality checks** (pmc.ncbi.nlm.nih.gov). The first five concern whether an alignment is reusable for system evaluation, while the remainder address completeness and non-trivialness (pmc.ncbi.nlm.nih.gov). Their discovery of **30 correspondences requiring removal** demonstrates that a reference alignment itself needs acceptance controls (pmc.ncbi.nlm.nih.gov). It does not establish a universal expected rejection percentage because the denominator and error mechanisms are corpus-specific.

Precision is the share of suggested correspondences that are correct, recall is the share of reference correspondences recovered, and F-measure is their harmonic mean (pmc.ncbi.nlm.nih.gov). These metrics evaluate sets against a reference. They do not by themselves test the consequences of a particular predicate for a particular consumer.

Semantic evaluation adds another caution. Standard precision and recall have limitations under OWL 2 semantics (link.springer.com). If merged ontologies and mappings are incoherent, published guidance recommends repair before semantic precision and recall are computed (link.springer.com). When a reference is incomplete, even semantic precision may be misleading, so only semantic recall may be appropriate (link.springer.com).

Human agreement is also evidence, not truth by default. One OAEI uncertain reference alignment derived confidence from agreement among **20 reviewers** (oaei.ontologymatching.org). Expert panels and consensus remain practical ways to construct references (www.cambridge.org), but acceptance should retain disagreement, rationale, and reviewer identity when they matter.

Review summary calculation

For an actual release, report mutually exclusive row outcomes with a common denominator:

- **Approval rate:** approved rows divided by all reviewed rows.
- **Downgrade rate:** downgraded rows divided by all reviewed rows.
- **Rejection rate:** rejected rows divided by all reviewed rows.
- **Deferral rate:** deferred rows divided by all reviewed rows.
- **Closure surprise rate:** newly inferred endpoint pairs that fail review divided by all newly inferred endpoint pairs reviewed.
- **Drift impact rate:** previously active rows changed after release review divided by all active rows touching changed entities.

Counts should accompany percentages, and deferred rows must remain in the denominator. No cited source establishes a universal pass-rate benchmark, so a report should not manufacture one. Trend comparisons are meaningful only when predicate mix, consumer scope, sampling method, and review policy remain comparable.

Implications and Future Directions

The acceptance method changes procurement and governance questions. A mapping product should be asked to export assertions, provenance, endpoint versions, and review state. A reasoning component should expose materialized consequences. A terminology service should provide resolvable identifiers and release history. A consumer should supply executable positive, negative, and boundary cases. No single score substitutes for those observable properties.

Persistent identity and semantic drift must be handled together. Getty's editorial guidance prioritizes keeping subject identifiers persistent during merges (www.getty.edu). Crossref similarly requires the new owner's DOI prefix in transfer updates (www.crossref.org). These practices preserve addressability through organizational or editorial change, but consumers still need to retest meaning and use.

Mapping governance will also need clearer confidence semantics. A numeric confidence can mean model output, annotator agreement, evidence strength, or calibrated probability. SSSOM allows a value between 0 and 1 (github.com), but a producer should state the estimator, reference population, calibration method, and intended decision threshold before consumers interpret it probabilistically.

Finally, mapping publication should become more explicit about status. GRC Schema's published rules provide a current example: notes identify endpoint versions, and a changed canonical mapping requires a vocabulary version change (grcschema.org). The general lesson is not to copy one rule set. It is to make acceptance policy, versions, and changes inspectable by consumers.

Frequently Asked Questions (FAQs)

What is an ontology crosswalk acceptance test?

It is a black-box review that decides whether a specific, versioned mapping claim is safe for named downstream uses. It verifies endpoint evidence, predicate semantics, inferred consequences, consumer queries, and version drift, then records approval, downgrade, rejection, or deferral.

What is the difference between exact match and close match?

In SKOS, both are symmetric. **Exact match is transitive** and expresses high-confidence interchangeability across a wide range of retrieval applications. **Close match is not transitive** and supports interchangeability only in some applications. Therefore, an A to B to C chain can imply A exactMatch C, but cannot imply A closeMatch C merely from two close mappings.

How should broader and narrower concept mappings be validated?

Validate direction first. For `broadMatch`, the subject is more specific and the object more generic. Then test roll-up, drill-down, exclusions, and aggregate counts. Materialize the inverse `narrowMatch` consequence and verify that neither direction is mistakenly consumed as equivalence.

Is SKOS exactMatch the same as owl:sameAs?

No. `owl:sameAs` states that two individuals are the same resource, while `skos:exactMatch` is a concept-mapping relation designed for cross-scheme interchangeability. `owl:equivalentClass` is different again because it asserts equal class extensions. The endpoint kinds and downstream logic determine which category is applicable.

Which ontology alignment evaluation metrics belong in acceptance?

Precision, recall, and F-measure describe alignment-set performance against a reference. Relation-aware and semantic metrics help when predicates and entailments matter. Row acceptance additionally needs consequence inspection and consumer-query tests because aggregate metrics can hide a small number of consequential false equivalences.

What should happen when no predicate fits?

Record **no acceptable mapping** with a reason and reviewed evidence. Do not leave a blank and do not force the weakest available relation if even `relatedMatch` would mislead consumers. An unmapped row can be the correct governed decision.

When should an accepted mapping be reviewed again?

Retest when either vocabulary publishes a release, an endpoint changes meaning, hierarchy, or status, a concept splits or merges, mapping policy changes, the consumer action expands, or an observed result violates the fixture. Also record a periodic backstop date for mappings whose sources do not publish machine-readable change events.

Conclusion

An ontology crosswalk acceptance test is a controlled decision about a semantic claim. It starts with the consumer action, fixes the source and target releases, selects the weakest adequate predicate, inspects the predicate's formal consequences, executes positive and negative boundary queries, and records a versioned outcome. This makes approval defensible and makes downgrade, rejection, deferment, and no acceptable mapping equally legitimate results.

The critical distinction is between **candidate plausibility** and **accepted consequence**. Labels and confidence scores can prioritize work, but the predicate determines what consumers may infer. Exact, close, broader, narrower, related, identity, and class-equivalence claims are not interchangeable. Their endpoint kinds, direction, symmetry, inverses, transitivity, and disjointness must be tested explicitly.

The final artifact should be reproducible: identifiers, observed labels and definitions, releases, predicate, justification, evidence, author and reviewer roles, fixture results, decision, observation date, and retest trigger. Metrics can summarize a release, but no universal pass rate replaces row-level evidence. A high-quality crosswalk is therefore not the one with the most links. It is the one in which every active link has survived the consequences that its actual consumers will rely on.

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/skos-reference
2. https://mapping-commons.github.io/sssom/1.0/spec-model
3. https://github.com/mapping-commons/sssom/blob/v1.0.0/src/sssom_schema/schema/sssom_schema.yaml
4. https://mapping-commons.github.io/rda-fair-mappings/recommendation
5. https://www.rd-alliance.org/groups/fair-mappings-wg/outputs
6. https://www.rd-alliance.org/recommendations-and-outputs/outputs
7. https://pmc.ncbi.nlm.nih.gov/articles/PMC3337267
8. https://pmc.ncbi.nlm.nih.gov/articles/PMC6971340
9. https://www.loc.gov/apis/additional-apis/linked-data-service
10. https://www.gov.uk/government/publications/open-standards-for-government/persistent-resolvable-identifiers
11. https://www.cambridge.org/core/journals/knowledge-engineering-review/article/user-validation-in-ontology-alignment-functional-assessment-and-impact/F68413CC283DAC0F3DF36C46000EC593
12. https://academic.oup.com/database/article/doi/10.1093/database/baac035/6591806
13. https://www.w3.org/TR/dwbp
14. https://op.europa.eu/en/web/eu-vocabularies/news/-/blogs/corporate-body-new-structure
15. https://www.ncbi.nlm.nih.gov/books/NBK9682?report=reader
16. https://www.nlm.nih.gov/research/umls/implementation_resources/query_diagrams/er7.html
17. https://op.europa.eu/documents/d/eu-vocabularies/immc-technicalities
18. https://www.w3.org/TR/skos-primer
19. https://www.dublincore.org/specifications/dublin-core/dcmi-terms/terms/creator
20. https://arxiv.org/abs/1310.4156
21. https://oaei.ontologymatching.org/2024/conference/index.html
22. https://oaei.ontologymatching.org/2025/beyondequivalence/Results.html
23. https://www.getty.edu/research/tools/vocabularies/contribute.html
24. https://www.nlm.nih.gov/pubs/techbull/so26/so26_umls_40.html
25. https://github.com/mapping-commons/sssom/issues/530
26. https://github.com/mapping-commons/sssom/releases/tag/v1.1.0a5
27. https://www.bioontology.org/wiki/BioPortal_Help
28. https://pmc.ncbi.nlm.nih.gov/articles/PMC5715990

29. https://link.springer.com/article/10.1186/s13326-017-0162-9
30. https://www.getty.edu/publications/vocabularies-editorial-guidelines/aat-guidelines/2_general_guidelines/2.2
31. https://www.crossref.org/documentation/register-maintain-records/maintaining-your-metadata/updating-your-metadata
32. https://grcschema.org/docs/reference/skos-mapping-decisions

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.