Measuring Object-Level Retrieval with Index Layers in Cultural Heritage Conservation Records

Andrea Gulli[1], Michele Geronazzo[2] and Massimiliano Guido[1]

[1]Department of Musicology and Cultural Heritage, University of Pavia, Corso Garibaldi 178, Cremona, 26100, Italy
[2]Department of Engineering and Management, University of Padova, Stradella S. Nicola 3, 36100, Vicenza, Italy

Abstract

In cultural heritage, conservation RDF often records evidence on object parts, while users query the described object. This paper presents a procedure for measuring how much of that evidence an application profile makes reachable from the object. The same object-level queries are run over a profile that preserves only fully-articulated component paths and over a profile that additionally exposes selected component-derived answers through a separate retrieval index; the index serves as a measurement instrument, holding the query form constant so the profile is the only variable. Applied to 158 Ligatus Linked Conservation Data board-reattachment records encoded in CIDOC CRM, across four query families covering modifications, conditions, treatment materials, and techniques, the procedure identifies 1737 URI-level answers that are documented in the source graphs but unreachable from the described object in 105 of the 158 records. Conditions and techniques are unreachable in every affected record, and roughly a third of recovered answers are supported two P46 links below the object. Each recovered answer is traced to its source component, CRM path, and index rule. The proposed method provides a repeatable way to quantify the object-level retrieval commitment implied by an application profile, without asserting new CRM modification or condition statements.

Keywords

CIDOC CRM, ontology evaluation, application profiles, retrieval indexing, conservation documentation, cultural heritage

1. Introduction

A historical musical instrument is rarely documented only as one undifferentiated object. The Theewes Claviorgan[1], held by the Victoria and Albert Museum, is a useful example: a 1579 London instrument combining organ and harpsichord structures, whose catalogue record describes extensive losses, repairs, replacements, overpainting and repainting, and a present condition in which the case and only fragments of the works survive. Such a record is not one statement about one object; it is a layered history of surviving parts, missing parts, altered surfaces, and changing material states. A user asking "what conservation work is documented for this instrument?" expects an answer at object level, even when the relevant evidence is documented more precisely at the level of parts, materials, techniques, or intervention events. The same layering governs other conservation domains. In book conservation, from which the evaluation corpus of this paper is drawn, treatment evidence is recorded against boards, textblocks, endleaves, and spine components rather than against the volume as a whole.

This is a practical application-profile problem for CIDOC CRM conservation data. CIDOC CRM can represent the described object, its components, modification activities, conditions, treatment materials, and techniques through explicit paths [1, 2]. That precision is valuable: a condition of the soundboard is not necessarily a condition of the whole instrument. At the same time, a query that starts from the object URI can miss evidence that is present in the graph but attached below a component path. The question

SWODCH'26: International Workshop on Semantic Web and Ontology Design for Cultural Heritage, Workshop days of ISWC 2026, Bari, Italy

andrea.gulli@unipv.it (A. Gulli); michele.geronazzo@unipd.it (M. Geronazzo); massimiliano.guido@unipv.it (M. Guido)
https://remake.unipv.it/persone/andrea-gulli/ (A. Gulli); https://www.unipd.it/contatti/michele.geronazzo (M. Geronazzo); https://remake.unipv.it/research-teams/ (M. Guido)
0000-0002-3000-1966 (A. Gulli); 0000-0002-0621-2704 (M. Geronazzo); 0000-0002-1398-7989 (M. Guido)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

[1]https://collections.vam.ac.uk/item/O60635/the-theewes-claviorgan-claviorgan-theewes-lodewyk.

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-09

CEUR-WS.org/Vol-4320/paper5.pdf

is therefore how much object-level retrieval changes when a profile preserves only fully-articulated component paths, or also provides a traceable retrieval index for selected component evidence.

This paper evaluates that question on the Ligatus Linked Conservation Data board-reattachment pilot [3]. It compares two application profiles over the same CIDOC CRM source evidence. The *fully-articulated path profile* preserves the source graph unchanged: component-derived evidence is reachable only by following the relevant component paths. The *retrieval-index profile* preserves those source paths but adds a separate index layer using custom predicates. These index links make selected component-derived answers retrievable from the described object without asserting new CIDOC CRM modification or condition statements. The access commitment therefore sits in the profile, as a set of dedicated link types, rather than in individual queries: the same standard query runs unchanged over both.

The contribution is a measurement procedure, not a proposed enrichment step. It runs the same object-level query families over both profiles, counts the answers made visible by the retrieval index, and preserves a trace from each recovered answer back to the component path and index rule that support it. The index layer is the instrument rather than the finding: as Section 4.1 makes explicit, a component-aware query using recursive property paths would reach the same modification answers. What is measured is therefore a property of the corpus—where conservation evidence sits relative to the object node—and what it yields is a quantified basis for the profile-design recommendation of Section 6.

2. Background

Ontology-based conservation documentation requires more than choosing a reference model. Historical and conservation records contain object structure, condition observations, interventions, treatment materials, techniques, removed or added parts, and changes of interpretation over time. Encoding such information requires standards, application profiles, modelling guidelines, and tools that make local choices explicit [4, 5, 6]. CIDOC CRM is widely used as a reference ontology for event-centred cultural heritage documentation [2, 1], but its generality means that implementations must still decide which paths are expected for a task and what retrieval support is offered to users.

Application profiles constrain a broad model for a project, data source, workflow, or retrieval task [7]. The present experiment uses the CIDOC CRM distinction between fully-articulated paths and shortcuts as motivation. In CIDOC CRM, a shortcut can provide a compact property for a longer path, while the fully-articulated path remains the more explicit representation. For the retrieval problem addressed here, however, asserting a CRM shortcut-like statement such as "the object was modified by this E11 activity" would change the meaning of the source evidence when the documented intervention concerns a component. The index predicates used here are instead custom retrieval links, leaving the component-level CRM statements intact.

The evaluation is task-based: it asks whether a profile supports object-level conservation queries, following the use of competency questions and task-based ontology evaluation [8, 9, 10]. The governing question is the one a user puts to a catalogue record—*CQ0: what conservation work is documented for this object?*—and the four query families used below are its refinements: which modification activities are documented (CQ1), which conditions are documented in treatment-relevant components (CQ2), and which materials (CQ3) and techniques (CQ4) those activities employed. CQ2 is deliberately narrower than asking for every condition recorded below the object, because the experiment focuses on conditions linked to conservation-treatment contexts. Whether CQ0 is answerable at all from the object URI is what the profile comparison measures.

Competency-question evaluation need not be binary: instance-based methods have used CQ answers to quantify alignment quality [11]. In cultural heritage, profiling work has measured how variably CIDOC CRM is instantiated across datasets [12], and recent question-answering work shows how far query construction depends on multi-hop path patterns [13]. None of these measures the retrieval cost of a declared profile: [11] scores a mapping, [12] characterises realised structure across corpora, and [13] compensates for that structure inside the query. Holding the source graph, information need,

answer type and query form fixed, and varying only the object-level access layer, turns the CQ verdict into a per-record, per-predicate count in which each additional answer stays traceable to its supporting source path.

3. Profile Comparison

3.1. CIDOC CRM Vocabulary for Conservation Retrieval

The proposed experiment uses only a small part of CIDOC CRM. The described conservation object is treated as an E22 Human-Made Object. Its physical parts, such as boards, covers, textblocks, and endleaves, are connected to it through P46 *is composed of*. In this paper, these parts are called components. Conservation interventions are represented as E11 Modification activities. E11 is the CIDOC CRM class for activities that modify a physical thing; in the Ligatus corpus, these activities are interpreted as conservation-treatment actions. A thing modified by an E11 activity is connected to that activity through P31i *was modified by*, the inverse of P31 *has modified*. Conditions are represented as E3 Condition State nodes attached to objects or components through P44 *has condition*. Treatment materials are read from E11 activities through P126 *employed*, and techniques through P32 *used general technique*.

Throughout this paper, the notation $o \rightarrow \mathrm{P46}^{+} \rightarrow c$ denotes a recursive path: $\mathrm{P46}^{+}$ follows the *is composed of* relation one or more times, reaching components at any depth below the described object. The *depth* of a component is the minimum number of P46 links separating it from the described object: one step is depth 1, a component reached through an intermediate component is depth 2, and so on. The depth columns in Table 1 record this value for each index-recovered answer.

3.2. Ligatus Evaluation Corpus

The evaluation uses 158 RDF records from the Ligatus Linked Conservation Data board-reattachment pilot [3]. Board reattachment is a common treatment in book conservation: the boards of a historic binding become detached from the textblock and are refixed, by approaches and materials that have varied considerably over the last half-century. The pilot gathered documentation of this one treatment from four institutions—the Bodleian Library, the Library of Congress, The National Archives (UK), and Stanford Libraries—converting free-text reports, spreadsheets, and rekeyed handwritten notes into CIDOC CRM RDF, with local terminology aligned to the Getty AAT and the Language of Bindings. Because the treatment is by definition an operation on a part, the records describe each binding as a structure of boards, covers, textblocks, endleaves, and spine components, and attach most conservation evidence to those parts rather than to the volume. Each RDF file is treated as one record; since the data are already supplied as CIDOC CRM RDF, the procedure reads the RDF structure directly.

The evaluation first builds a profile-independent record table. For each RDF file it selects the described object structurally: it collects all E22 Human-Made Object nodes, computes for each the number of distinct nodes reachable through transitive P46 *is composed of* links, and takes the E22 node with the largest component closure, that being the entity from which all component-level evidence is expected to be retrieved. Ties are broken first by preferring candidates with production evidence through P108i *was produced by*, then by lexical URI order. Record labels and record-specific identifiers are not used as selection evidence. The largest-closure rule selected the described object without tie-breaking in 157 of the 158 records; the single tie, between three E22 candidates with equal-sized P46 closures, was resolved by lexical URI order. For the selected object, the table records its components, the minimum number of P46 steps needed to reach each component, the typed E11 Modification activities linked through P31i, and the values relevant to the four query families: P44 conditions, P126 treatment materials, and P32 techniques. This extraction is performed before the two profile answer sets are compared.

The corpus is strongly component-oriented. Across the 158 records, the extraction finds 977 E11 Modification activities. Typed E11/P31i links occur on nodes other than the selected object in 150 records and directly on the selected object in 106 records. Eligible P46-closure component targets occur

in 105 records. A component-faithful representation can therefore be semantically precise while leaving a direct object-level query blind to most of it.

3.3. Retrieval-Index Profile

The *retrieval-index profile* keeps the same source graph and adds a separate index graph using the custom predicates `id:indexModification` and `id:indexCondition`. Let o be the selected described object, c a component reachable from o by one or more P46 links, m a typed E11 Modification activity, and q a condition state. Two index rules generate the index triples.

E11 index rule. For each source path $o \rightarrow \text{P46}^{+} \rightarrow c \rightarrow \text{P31i} \rightarrow m$, with m explicitly typed as `E11 Modification`, the rule adds $o \rightarrow$ `id:indexModification` $\rightarrow m$. Materials (P126) and techniques (P32) are not copied into the index; they remain on the source E11 activity and become object-retrievable through the indexed activity.

Condition index rule. For each source path that satisfies the E11 index rule conditions and additionally has $c \rightarrow \text{P44} \rightarrow q$, the rule adds $o \rightarrow$ `id:indexCondition` $\rightarrow q$. Requiring the E11 co-occurrence restricts lifting to conditions documented in conservation-treatment contexts, preventing ambient component observations—condition states recorded independently of any intervention—from being surfaced at object level.

For each generated triple, the trace records the source component c, the supporting E11 activity m, the rule applied, the generated triple, and the minimum P46 depth of the path from o to c. If an answer has several supporting paths, all are retained; depth summaries use the shortest.

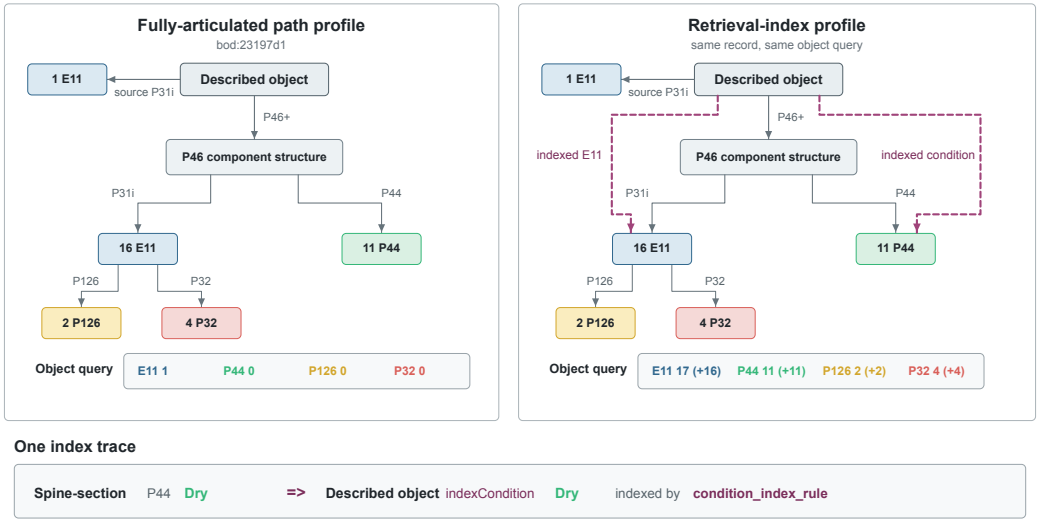


Figure 1: Illustrative record-level comparison. The same object-level query is run over both profiles; the retrieval-index profile recovers component-derived conservation evidence through `id:indexModification` and `id:indexCondition`. The trace example reports that the condition labelled "Dry", attached by P44 to the component labelled "Spine-section", supports the generated triple described object `id:indexCondition` "Dry" via the condition-index rule.

One record. Figure 1 shows both rules on a single record of the Bodleian subset, `bod:23197d1`. Its evidence is distributed rather than absent: one E11 Modification hangs from the described object, while sixteen more sit on components reached through P46, together with eleven conditions, two materials, and four techniques. The object-level query returns that one modification and nothing else, although the rest is present in the same file. The E11 index rule then links the sixteen component activities to the object—making their P126 materials and P32 techniques reachable in turn—and the condition index rule links the eleven treatment-context conditions, taking the same query to 17, 11, 2 and 4 answers. One line of the trace records why: the condition "Dry" is attached by P44 to the component "Spine-section",

which carries a typed E11 activity inside the object's P46 closure, so the rule generates *described object* `id:indexCondition` "Dry".

4. Measuring Object-Level Accessibility

4.1. Object-Level Query Families

For each record, the evaluation runs the four query families of CQ1–CQ4 over both profiles, each rooted at the described object. The answer families are the same in both profiles, but the available object-level access paths differ. In the fully-articulated path profile, the query uses only CIDOC CRM properties asserted in the source graph: P31i for E11 Modification activities, P44 for conditions, and P126 or P32 from the retrieved E11 activities. In the retrieval-index profile, the source graph is unchanged, but two additional predicates are available: `id:indexModification` for indexed E11 activities and `id:indexCondition` for indexed conditions. Materials and techniques are still read from source E11 activities through CIDOC CRM P126 and P32; they are not copied into the index graph.

A component-aware SPARQL query using recursive property paths—for instance $\texttt{P46}^{+}\texttt{/P31i}$ to reach component-level E11 activities—would recover the same modification answers as the E11 index rule for that family. The evaluation uses the same non-recursive, object-rooted query template over both profiles. This isolates the profile as the variable under test. Automatic path-pattern query generation [13] is an alternative instrument for the same gap: it builds the traversal into each generated query, where a declared profile builds it in once. Either closes the gap; the purpose here is to measure it.

Candidate modification activities are counted only when explicitly typed as E11 Modification.

4.2. Answer-Set Coverage

For each record R and answer family F, the evaluation stores a normalised URI-level answer set per profile: $A_{\text{full}}(R, F)$ from the fully-articulated path profile and $A_{\text{index}}(R, F)$ from the retrieval-index profile. The central measure is their set difference, $IR(R, F) = A_{\text{index}}(R, F) \setminus A_{\text{full}}(R, F)$. $IR(R, F)$ is the set of *index-recovered answers*.

The comparison uses URI-level sets. Duplicate triples do not create duplicate answers, labels are not used to merge answers, and the same URI appearing in two records is counted once per record. The index profile never removes source answers, so the answer-set difference measures additional object-level accessibility, not replacement.

4.3. Visibility Diagnostics

The evaluation also reports *direct-access coverage* for each answer family: the per-record distinct-URI answer counts are summed, and the fully-articulated profile total is divided by the retrieval-index profile total. This quantity is deliberately not called recall. Its denominator is the retrieval-index profile's own output rather than the documented answer space, so the index profile scores 100% on it by construction; the outside-path counts introduced below report answers that neither profile reaches. A low value means that many retrieval-index-profile answers were not directly reachable from the described object in the fully-articulated path profile. This differs from RDF query completeness, which asks whether a source contains sufficient information to guarantee complete query answers [14]; here both profiles operate over the same fixed source graph.

The evaluation also reports *dark records*. A record is dark for family F when $|A_{\text{full}}(R, F)| = 0$ while $|A_{\text{index}}(R, F)| > 0$. Dark records matter because a direct object-level query then gives no indication that relevant documented evidence exists anywhere in the record.

The procedure also establishes the documented answer space in the source data, so the retrieval-index profile is not mistaken for complete retrieval of the RDF graph. For each record it counts all explicitly typed E11 Modification nodes and all objects of P44, P126, and P32—the latter predicate-based, not

filtered by range typing or by connection to the selected object—and divides them into answers visible without the index layer, index-recovered answers, and answers reached by neither profile.

Three supporting diagnostics explain the differences in object-level accessibility. Conditions on components without a P31i link to a typed E11 Modification fall outside the condition index rule. Evidence attached to components outside the selected object's transitive P46 closure falls outside both index rules. Records containing no eligible component evidence remain unchanged across profiles. The depth diagnostic addresses a separate issue: answers supported below depth 1 require recursive P46 traversal or an explicit index and cannot be recovered by a single-step component query.

5. Results

5.1. Index-Recovered Answers

Coverage. Across the full corpus, the retrieval-index profile recovers 1737 URI-level answers that are not returned by the same object-level queries in the fully-articulated path profile. The effect changes 105 of 158 records. E11 activities account for 43.5% of those answers and techniques for a further 26.4%, against 24.9% for conditions and 5.2% for materials (Figure 2). Index dependence is strongest for E11 activities and techniques, each affecting 105 records, substantial for conditions (95 records), and much weaker for materials (24 records). The profile effect is therefore not uniform across CIDOC CRM predicates.

Figure 2 shows that the effect is broad but heterogeneous. Of the 53 unchanged records, 52 contain no eligible indexed content under the implemented rules and one is already object-visible; none is explained by component evidence duplicating visible URIs.

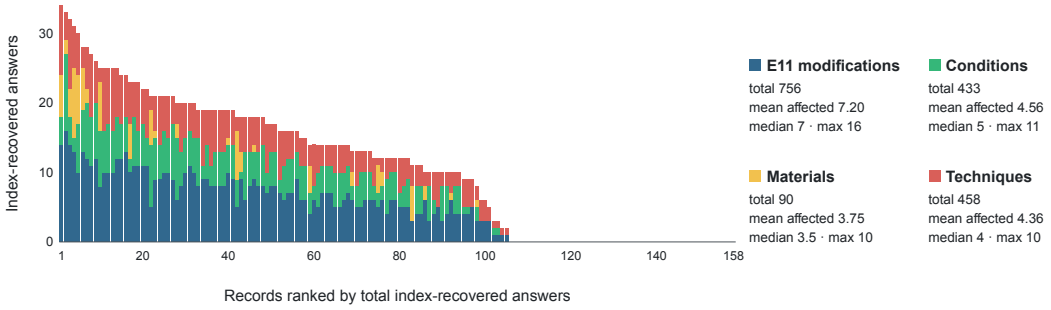


Figure 2: Record-level composition of index-recovered answers across all 158 records. Each stacked bar is one record, ranked by total recovered answers. The empty right-hand tail corresponds to the 53 records whose answer sets do not change.

5.2. Object-Level Visibility

Table 1
Object-level visibility diagnostics. *Direct access* is the share of retrieval-index-profile answers already reachable in the fully-articulated path profile; it is not recall against the documented answer space. *Dark / non-empty* counts records where the fully-articulated profile returns nothing, out of those where the retrieval-index profile returns at least one answer. Depth counts refer only to index-recovered answers.

Answer family	Full	Index	Direct access	Dark / non-empty	Outside paths	Depth 1	Depth 2
E11 modifications	117	873	13.4%	0/106	104	493 (65.2%)	263 (34.8%)
Conditions	0	433	0.0%	95/95	401	312 (72.1%)	121 (27.9%)
Materials	467	557	83.8%	21/97	245	41 (45.6%)	49 (54.4%)
Techniques	0	458	0.0%	105/105	185	288 (62.9%)	170 (37.1%)

Table 1 reports the visibility diagnostics. Conditions and techniques are completely dark along the implemented object-level paths. E11 activities behave differently: their direct-access coverage is low, at 13.4%, but no record is E11-dark, so a direct object-level query always returns something for that family—it simply returns a small fraction of what is documented. Materials have comparatively high direct-access coverage but still include 21 dark records. The material asymmetry reflects the Ligatus encoding: many P126 values hang from E11 activities already reachable from the object, whereas P32 values become visible only through indexed component-level activities. Two of the zeros in Table 1 are worth stating plainly: no P44 condition is attached to a described object anywhere in the corpus, and none of the 117 object-visible E11 activities carries a P32 technique, although those same activities account for 467 P126 material values. The object node is therefore not an impoverished view of the component record but a different register of statement, and the description of how a binding was treated sits entirely below P46.

The outside-path counts limit the claim. The retrieval-index profile improves object-level accessibility, but it is not complete retrieval of every documented conservation statement in the RDF. This is especially visible for conditions: 401 condition answers are documented in the source graphs but remain outside the implemented object-level paths. Most of these are excluded by the condition-index constraint: 320 conditions are attached to P46 components that lack a typed E11/P31i link, and 81 are attached to components outside the selected P46 closure.

The depth counts show that indexing is not only compensating for one-step component paths: 34.8% of recovered E11 activities and 27.9% of recovered conditions are supported at depth 2. In this corpus, depth 2 evidence is most often associated with textblocks, endleaves, spine sections, and left or right sections nested beneath another P46 component.

6. Discussion and Conclusion

The central issue behind these numbers is a tradeoff between structural precision and object-level accessibility. Fully-articulated component paths preserve where the evidence was documented: an intervention on a component should remain identifiable as such. At the same time, that precision can make ordinary object-level discovery incomplete.

The retrieval-index profile addresses the accessibility side of this tradeoff without changing the CIDOC CRM source assertions: it is a traceable access layer over selected source paths, not a claim that component modifications or component conditions are literally whole-object statements.

Limitations. First, Ligatus is a board-reattachment pilot, not a representative sample of all conservation documentation. Second, the index rules are not corpus-agnostic, and this bounds what the numbers transfer. CIDOC CRM's expressive latitude means the same conservation fact can be encoded through different property chains in different projects: a treatment reached through P31i in one dataset may be modelled as a broader E7 Activity in another, and a condition may hang from the component, from a separate observation event, or from an assessment record held apart from the treatment. The rules used here were derived by inspecting the conventions of a single pilot, so porting them means re-deriving them against the target corpus rather than re-running them unchanged. The measurement procedure does port: what it asks of a new corpus is a declared object-selection rule, a declared set of index rules, and the same held-constant query form. The counts in Table 1 are therefore specific to Ligatus, while the method that produces them is not. Third, the evaluation covers only the implemented conservation paths: E11 Modification, P44 condition, P126 treatment material, and P32 technique. Fourth, the counts are URI-level answer counts: they are reproducible and traceable, but they are not a complete conceptual analysis of all conservation statements in the graphs.

Conclusions. This paper set out a measurement procedure for object-level answer coverage: hold the query form constant, vary only the access layer, count the difference per record and answer family, and keep every counted answer traceable to the CRM path supporting it. Applied to the Ligatus board-reattachment corpus, it finds 1737 documented answers that are not reachable from the described object—every condition and every technique, 87% of modification activities, and 16% of treatment

materials—with a third of them lying two P46 links down.

That measurement motivates a design consequence rather than asserting one independently. Application profiles for component-rich CIDOC CRM datasets should declare their object-level retrieval commitment explicitly: if object-level discovery is expected to include component evidence, the profile should either provide a transparent retrieval index or mandate equivalent component-aware queries. Without such a declaration a missing object-level answer is ambiguous—the evidence may be absent, it may exist only on a component path, or the query may not be using the profile's intended access layer—and the present measurements show that in a real corpus the second case dominates.

Future work. Applying the procedure to a further Linked Conservation Data pilot would separate findings specific to this corpus from properties of component-rich CRM encoding in general, and would test whether index rules can be derived semi-automatically from observed encoding patterns rather than by hand. Extending beyond the four implemented paths—in particular to the conditions outside treatment contexts that account for 320 of the 401 outside-path condition answers—would establish whether the gap measured here is characteristic of the data or an artefact of the paths selected.

Data and Code Availability

Code, profile outputs, index rules, comparison tables, traceability ledgers, and diagnostic summaries are available at https://github.com/gullogullo/crm-modelling-choice-benchmark.

Acknowledgments

This research is funded by REM@KE Synergy ERC Research Project 101167101. The authors thank the Ligatus and Linked Conservation Data contributors for making the board-reattachment pilot data available, and Jacopo Uberti, Matteo Lombardi, and Sebastian Kirsch for discussions on musical instrument collection records and conservation problems.

Declaration on Generative AI

During the preparation of this work, the authors used GPT-5.5 and Claude Sonnet 4.6 for text drafting, summarisation, literature search support, and code drafting. After using these services, the authors reviewed, verified, and edited the content as needed, and take full responsibility for the publication's content.

References

[1] CIDOC CRM Special Interest Group, CIDOC Conceptual Reference Model, n.d. URL: https://www.cidoc-crm.org/, accessed 2026-06-22.

[2] M. Doerr, The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata, AI Magazine 24 (2003) 75–92. doi:10.1609/aimag.v24i3.1720.

[3] A. Velios, Datasets for the Pilot of Linked Conservation Data Project, 2022. doi:10.25441/arts.18551366, dataset used for the Linked Conservation Data board-reattachment pilot.

[4] R. Lieu, A. Campagnolo, Modelling Linked Data for Conservation: A Call for New Standards, KULA: Knowledge Creation, Dissemination, and Preservation Studies 6 (2022) 1–8. doi:10.18357/kula.232.

[5] E. Moraitou, Y. Christodoulou, G. Caridakis, Semantic Models and Services for Conservation and Restoration of Cultural Heritage: A Comprehensive Survey, Semantic Web 14 (2023) 261–291. doi:10.3233/SW-223105.

[6] A. Gulli, M. Geronazzo, M. Guido, DIMA: Toward a Multimedia Inspection Workbench for Musical Cultural Heritage Records, in: 2026 International Conference on Content-Based Multimedia Indexing (CBMI), IEEE, 2026. Forthcoming.

[7] S. Barzaghi, I. Heibi, A. Moretti, S. Peroni, Developing Application Profiles for Enhancing Data and Workflows in Cultural Heritage Digitisation Processes, in: International Semantic Web Conference, Springer, 2024, pp. 197–217. doi:10.1007/978-3-031-77847-6_11.

[8] M. Grüninger, M. S. Fox, The Role of Competency Questions in Enterprise Engineering, in: A. Rolstadås (Ed.), Benchmarking — Theory and Practice, IFIP Advances in Information and Communication Technology, Springer, Boston, MA, 1995, pp. 22–31. doi:10.1007/978-0-387-34847-6_3.

[9] C. Bezerra, F. Freitas, Verifying Description Logic Ontologies Based on Competency Questions and Unit Testing, in: M. Abel, S. R. Fiorini, C. Pessanha (Eds.), Proceedings of the X Seminar on Ontology Research in Brazil and I Doctoral and Masters Consortium on Ontologies, volume 1908 of *CEUR Workshop Proceedings*, 2017, pp. 159–164. URL: https://ceur-ws.org/Vol-1908/paper16.pdf.

[10] R. Porzel, R. Malaka, A Task-based Approach for Ontology Evaluation, in: ECAI Workshop on Ontology Learning and Population, Valencia, Spain, volume 1, Citeseer Valencia, Spain, 2004.

[11] E. Thiéblin, O. Haemmerlé, C. Trojahn, Automatic Evaluation of Complex Alignments: An Instance-Based Approach, Semantic Web 12 (2021) 767–787. doi:10.3233/SW-210437.

[12] M. Mountantonakis, I. Theocharakis, Y. Tzitzikas, Why We Need Ontology-Specific Data Portals: A Case Study for CIDOC-CRM, in: Proceedings of the International Workshop on Semantic Web and Ontology Design for Cultural Heritage (SWODCH 2023), volume 3540 of *CEUR Workshop Proceedings*, 2023. URL: https://ceur-ws.org/Vol-3540/paper9.pdf.

[13] M. Mountantonakis, Y. Tzitzikas, Generating SPARQL Queries over CIDOC-CRM Using a Two-Stage Ontology Path Patterns Method in LLM Prompts, ACM Journal on Computing and Cultural Heritage 18 (2025) 12:1–12:20. doi:10.1145/3708326.

[14] F. Darari, W. Nutt, G. Pirrò, S. Razniewski, Completeness Management for RDF Data Sources, ACM Transactions on the Web 12 (2018) 18:1–18:53. doi:10.1145/3196248.