

# Computational Approaches for Addressing Problematic Terminology in Museum Catalogues: A Knowledge Graph of Museum Critical Cataloguing Guidelines

Erin Canning<sup>1</sup>[0000-0003-3383-272X]

<sup>1</sup> University of Oxford, Oxford OX1 2JD, UK  
erin.canning@eng.ox.ac.uk

**Abstract.** Many museums are interested in addressing the presence of problematic terminology in their catalogue data but are uncertain of what they should look for and what do to when a problematic term is encountered. There is an opportunity to address the lack of guidance and scarcity of resources available for museums looking to engage with this work through the use of linked open data in order to create a knowledge graph that links together machine-readable versions of terminology guidance from multiple sources and then further connects them to additional lexicographical resources. The resulting knowledge graph can then be leveraged to support critical cataloguing tasks—namely confirming a term as problematic in a given context and making a decision of how to proceed—by providing cataloguers with the information they need to perform their work. This approach seeks to meet the needs of museum professionals to have access to domain best practice while allowing for localization of recommendations. This research seeks to understand what kinds of language museums are concerned with, how museums are thinking about what makes a term problematic, what possible reparative actions museums feel are appropriate given different contexts, and what kinds of information museum professionals look for to make a decision about which action to take. This work is taking place as part of a Collaborative Doctoral Partnership studentship co-supervised by the University of Oxford and the Victoria and Albert Museum, and as such the Victoria and Albert Museum serves as the case study for the exploration of this research area.

**Keywords:** ontology, knowledge graph, linked open data, cultural heritage, problematic terminology

## 1 Introduction and Motivation

The Victoria and Albert Museum (V&A) first opened its doors to the public as the South Kensington Museum in 1857 and today contains an extraordinarily diverse collection of close to 1.7 million works of art and design objects. Just as diverse as the objects are the catalogue records that accompany them: this catalogue data contains evidence of the colonial history of the museum and of the relationship between Britain and the areas of the world in which Britain had colonial or proto-colonial interests, as objects from these parts of the world make up a significant portion of the museum's

collection. This evidence can appear in various forms, including the presence of “problematic terminology”: harmful and offensive language, colonial naming practices, euphemisms, clichés, and more. This issue is not unique to the V&A, and the museum field as a whole is wrestling with how to best address the problematics in their data.

One of the challenges for museums in addressing problematic terminology in catalogue data, also called critical cataloguing, is in determining where to begin: many museums may be interested in addressing issues in their data but are uncertain of what they should look for and what do to when a problematic term is encountered. The lack of guidance and scarcity of resources available for museums looking to engage with this work is well known [1–3]. There is an opportunity to address this through the use of linked open data (LOD): museum professionals can benefit from having a structured resource that makes clear how different terms are understood by different institutions, and how different museums are approaching critical cataloguing work. The use of LOD will enable the creation of a knowledge graph that can link together machine-readable versions of terminology guidance from multiple sources and then further connects them to additional lexicographical resources. The development of this resource will contribute to the understanding of how museums are thinking about what makes a term problematic and what reparative actions are possible given different contexts, as well as supporting museum workers engaged in this work by addressing the need to develop sector-wide guidance and support cross-institutional knowledge exchange.

## 2 Related Work

### 2.1 Addressing Problematic Terminology in Museum Catalogue Data

The Museums Association Decolonization Working Group produced a decolonization support guide in 2021, two pages of which are dedicated to collections cataloguing [3]. Resources aimed at the museum community have also been developed by individual people and institutions; for example, [4] includes a glossary of problematic terminology and suggestions for what other museums could do when encountering the terms in their own catalogue data. In the UK, [5] published a set of recommendations for conducting critical cataloguing work and [6] produced a guide for critical cataloguing in archives building on those recommendations. There is also now [7], a collaborative resource intended to “promote cross-institutional collaboration on inclusive description issues” [1]. Although these resources started locally, they respond to a field-wide need and seek to contribute to a collective effort.

Individual institutions are working to put these guidelines into action, and museums are forming internal terminology and cataloguing review groups to audit and remediate their catalogue data. Auditing involves reviewing catalogue data for instances of harmful language, such as specific words, or for colonial, dehumanizing, or otherwise problematic framings. Once problematic terminology has been identified it can be addressed in a number of ways; there is no single rule that would be appropriate in all cases and it is essential to take the context and purpose of the work into consideration [2, 8]. When documented locally, these lists of problematic terms and suggested reparative actions become terminology guidance documents for the use of cataloguers within a museum.

## 2.2 Ontologies for Addressing Problematic Terminology in Cultural Heritage

While numerous lists of offensive, sensitive, and problematic terminology exist, the most relevant related work is in ontologies for the representation of terminology and cultural heritage data. This is because the focus of this project is not producing a list of problematic language, but on enabling both an understanding of what ‘problematic terminology’ means in the context of cultural heritage data, and facilitating cross-institutional sharing of information on terms and suggested reparative actions. Therefore, the focus of this work is not on the terms themselves, but instead on how they are framed and considered by cultural heritage institutions, as evidenced by the structure and contents of terminology guidance documents. The most relevant existing ontologies are:

- CIDOC Conceptual Reference Model (CIDOC CRM): an ontology for the representation of cultural heritage data [9]. It is a stable standard widely used in the cultural heritage community and has been an ISO standard since 2006 (and updated in 2023).
- Ontology Lexicon (OntoLex): an ontology for the representation of linguistic data, namely lexica and dictionaries. It is a stable standard, and the most widely used model for publishing these linguistic resources as linked open data [10].
- Simple Knowledge Organization Schema (SKOS): an ontology for the representation of thesauri and terminology lists. It is a stable standard and is widely used for publishing controlled vocabularies as linked open data. It has been used to create LOD critical cataloguing vocabulary resources such as Homosaurus [11].
- Cultural Contexts Concept Scheme for Contentious Terminology (CULCO): an ontology developed in 2022 as part of the ‘‘Culturally Aware AI’’ project to describe [6]’s glossary section [12].

These ontologies do not exist in isolation: points of connection link them together into a base that can then be further developed (see section 6 below). CULCO, OntoLex and CIDOC CRM all assert relationships to SKOS, making it key for bringing them together. While these connections are usually encoded as inheritance relationships, others are more informally proposed in scope notes or wider documentation. Alignment between OntoLex and CIDOC CRM has also been previously proposed [13].

## 3 Problem Statement and Contributions

### 3.1 Research Questions

This project seeks to examine the high-level question of if and how computational approaches can support the goals of, and work involved in, critical cataloguing. In pursuit of knowledge towards this goal, this research aims to answer the following questions:

**RQ 1.** How can domain knowledge about the identification and remediation of problematic terminology in museum catalogues be structured as a data schema that is actionable for use in a technical system?

This research question addresses the hypotheses that textual terminology guidance documents produced by museums can serve as knowledge sources for building an ontology and actionable knowledge graph of the informational needs of reparative critical cataloguing practitioners, that a linked open data solution that makes use of an ontology for the representation of terminology guidance documents will promote inter-institutional collaboration and knowledge sharing in this area in a way that also allows for each location to retain their local specificity, and that combining the terms listed in terminology guidance documents with additional information from linked open data resources will enable us to understand what kinds of language museums consider to be potentially problematic.

**RQ 2.** How can museum records containing potential indicators of the presence of biased language be identified using existing computational methods?

This research question addresses the hypothesis that providing access to information about the historical and contemporary meanings of terms, along with information about the different reparative methods suggested by the museum field, will enable critical cataloguing practitioners to review records and take remediative actions with greater confidence and efficiency.

**RQ 3.** Can computational approaches enable critical engagement with catalogue data, and if so, how can they be developed? What might a methodology for the development and application of computational methods informed by the theoretical foundations and concerns of critical cataloguing look like?

This research question addresses the hypotheses that computational approaches can enable critical engagement with catalogue data when they are developed not to address only specific existing work practices but to meet the goals of critical cataloguing work as a movement in the museum field, and that the theoretical foundations and concerns of critical cataloguing can inform the methodology for developing computational tools and methods for working critically with catalogue data.

### 3.2 Outputs and Contributions

This project will produce three distinct yet connected technical outputs: an ontology for representing problematic terminology and suggested actions, a populated and linked knowledge graph that incorporates multiple terminology guidance documents and additional linked open data vocabularies, and a computational environment that supports target knowledge graph usage by enabling the exploration of museum catalogue data through the components of the knowledge graph: the presence of identified terminology alongside the display of information about the term and suggested remediative actions.

The contributions to knowledge that this project seeks to make include:

- Understanding what kinds of problematics museums are concerned with, evidenced by what kinds of information make up terminology guidance documents.
- Comparing between institutions what is being looked for, how it is being identified, and what is being recommended when instances are encountered.

- Demonstrating how computational approaches can aid in identifying catalogue records that have the potential to contain bias.
- Demonstrating how concepts from critical theory can inform methodology for data engineering projects situated in this domain.

## 4 Research Methodology and Approach

The initial modeling work for the design of an ontology for the representation of problematic terminology guidelines (described below in section 6) was undertaken using the Victoria and Albert Museum Terminology Guidance Document as the source document for analysis. Knowledge was elicited from the structure and content of this document, and was used to enumerate the requirements of the schema. These requirements were then compared to the affordances of existing ontologies, a gap was identified, and a modelling solution was proposed. The design approach taken for this intends to minimize the creation of new classes and instead focus on ways to further bring together existing LOD resources and the communities involved in their development, therefore bringing this project into conversation with the data communities most relevant to the work being undertaken on this project and contributing towards the reuse of existing resources in the linked open data domain. The design approach builds off of two ontology engineering approaches: the approach to designing and adopting ontologies for use in domain of cultural data from a feminist methodological perspective discussed in [14], and the method of lean ontology development [15]. Feminist theory is one of the fields of critical theory that directly informs critical cataloguing [16], and therefore this methodological approach is highly relevant for this project. The lean ontology development method discussed in [15] has also been trialed in the setting of a feminist linked open data project which validates the suitability of this approach for this project. These approaches will be used as the foundations for exploring an ontology engineering methodology informed by the theoretical foundations and concerns of critical cataloguing.

This model will be validated through review with museum staff, as well as through the integration of two additional terminology guideline documents relevant to the field: the glossary section of [4] and the contents of [7]. These are frequently cited resources: [4] is regularly looked to as a guiding document in this field, and [7] is crowd-sourced and expansive. Integrating these sources will validate the proposed schema as applicable beyond the single context, and comparisons of terms and suggestions between different sources could be analyzed as a form of inter-researcher or inter-institutional agreement. Linking this knowledge graph to other LOD resources will involve looking for term matches in the Getty Art & Architecture Thesaurus (a commonly used controlled vocabulary) and the Homosaurus vocabulary (a community-developed LOD vocabulary that is already working in the space of critical cataloguing) to assert relationships to terms published by those sources. This will work to expose how problematic terminology may be introduced into catalogue data through the use of externally-sourced vocabularies and highlight community-developed vocabularies as alternatives.

This project also seeks to incorporate the perspectives of museum professionals engaged in critical cataloguing by means of an ongoing set of interviews which have been

taking place since December 2023. Findings from these interviews about the requirements, pain points, and goals of critical cataloguers inform project developments.

## 5 Evaluation Plan

Evaluation will be done in three ways: through review with V&A staff, through the integration of additional guidelines, and through the use of competency questions over the final knowledge graph. The initial validation involved reviewing the conceptual model with key stakeholders from the V&A staff. These are staff who are responsible for creating and maintaining the museum's Terminology Guidance Document (TGD), and therefore are best positioned to evaluate whether the conceptual model matches their understanding of the document and domain. A second validation review was then carried out with the same staff to confirm that the suggestions had been accurately translated into the model. This was an important area for confirmation as the "suggestion" section is the core purpose of the TGD, and accuracy is therefore critical for the validity of the knowledge graph. Integrating [4] and [7] will confirm the interoperability of the structure of the ontology. This validation test will determine whether the kinds of information found in the TGD are the same kinds of information as found in other terminology guidelines, or if there are additional kinds of knowledge and relationships that will need to be accommodated, thus requiring revision of the ontology. This validation is being performed after the reviews with V&A staff to ensure that the base model being compared is a stable foundation for this stage of the work. Validation through competency questions will be performed once the additional terminology guidelines have been integrated and links to the additional vocabularies sources have been made. These questions will be developed in collaboration with the V&A stakeholders involved in the earlier validation work, and will be formulated to answer questions that they had when creating the original document and when engaging in catalogue review.

Evaluation of the computational environment will return to review with V&A staff. This evaluation will be intended to consider whether the use of the information contained in the knowledge graph is indeed useful for the practical or conceptual work involved in critical cataloguing: namely, if it enables cataloguers to review records and take remediative actions with greater confidence and efficiency, and if working with catalogue records through this lens supports a critical engagement with catalogue data. This evaluation step is closely related to the competency question validation step, but will move beyond it to include the experiences of stakeholders who would potentially use the outputs of this project in their own work.

## 6 Results

### 6.1 Conceptual Model of the V&A Terminology Guidance Document

The V&A holds regular cross-departmental meetings to discuss terminology questions and concerns, with the goal of coming to recommendations about how to address general and specific cases. Decisions are documented and inform the V&A Terminology

Guidance Document (TGD): a living document, developed in collaboration with the Interpretation Department and additional staff-led internal advisory groups, that is intended to support staff in making decisions of how to proceed when they encounter problematic terminology in catalogue records. The TGD contains three main sections:

1. Potentially problematic terms: in addition to a central term, the document lists common variants and forms that a cataloguer might encounter in object records.
2. Term description: a number of different ways that the term has been used or understood across different times and places, reminiscent of a diachronic dictionary.
3. Suggestions: a list of suggestions that could be considered based on specifics of the term encounter, such as who or what is being described, the original intended use of the term in the record, and the field in the record where the term is found.

The conceptual modelling stage of work separates each of these elements into separate entities and defines the relationships between them. There is also information about the TGD itself: as a living document, both the guideline and individual entries are written, edited, and informed by the expertise of various people over time. Therefore, the document itself as an information object is a necessary element of the model.

## 6.2 Gap Analysis of Existing Ontologies

The next stage of this work was to review existing linked open data ontologies in order to evaluate whether or not a new ontology is required to represent this information. CIDOC CRM, OntoLex, SKOS, and CULCO covered many of the periphery elements found in the TGD conceptual model, but lacked in core coverage as well as actions involved in the work of the domain. However, due to their existing formal and informal relationships, this set of ontologies provide a robust foundation to build from.

CULCO represents the glossary of [4]: this resource lists terms along with descriptions of their history, use, and possible sensitivities, as well as suggestions of actions for cataloguers to take when they encounter in the term in a museum record. The CULCO model is comprised of three classes and six properties. Individual terms are classed as *ContentiousIssues*, which have labels (*skos:XL:Label*)—the way that they are written—and *Suggestions*. This *Suggestion* class is the closest existing class for a suggestion as it is understood in the context of the TGD, but it is too lightweight to be used for an accurate translation of the conceptual model of the TGD as it does not expand on the concept of labels, to which it is directly connected, beyond using SKOS-XL directly. This is similar to how Homosaurus uses *skos:altLabel* for their “use for” concept [11]. As such, it was found to be a strong foundation but not sufficient on its own.

OntoLex, conversely, introduces a level of complexity in modelling lexical information beyond what is required for the representation of the TGD: using OntoLex to model term forms and descriptions would require the use of classes such as *Lexical-Concept* that, for the purposes of this context, would be purely intermediary and beyond the requirements. Similarly, CIDOC CRM introduces a level of complexity in modelling information about cultural heritage objects that exceeds the requirements of the TGD. Even with these added levels of complexity, the core elements elicited during the conceptual modelling stage are not sufficiently covered by existing ontologies.

### 6.3 Proposed Ontology

The ontology proposed at this stage addresses this gap by declaring six new classes and ten new properties—labelled with the prefix “caapt” for “Computational Approaches for Addressing Problematic Terminology”—as subclasses and subproperties of elements from one or more of these four ontologies. The new classes represent core concepts for the ontology: the types of authority documents being detailed in this space (*caapt:Guide* and subclasses *caapt:TerminologyGuide* and *caapt:StyleGuide*), the terms described in terminology guidance documents (*caapt:TermRoot*), the contexts of term uses (*caapt:UseContext*), and the suggestion instructions (*caapt:Suggestion*).

The ten proposed properties can be grouped by their constraints and superproperties:

- Properties which describe the details of use contexts: *caapt:used\_where*, *caapt:used\_when*, *caapt:used\_why*, *caapt:about\_who*. These take *caapt:UseContext* as domain, different classes as ranges, and are subproperties of *ontolex:usage*.
- Properties which describe relationships between different suggestions: *caapt:use\_along\_with*, *caapt:preferred*, *caapt:if\_not\_possible\_use*. These take *caapt:Suggestion* as domain and range, and are subproperties of *crm:P69\_has\_association\_with*.
- Properties which describe relationships between terms and suggestions: *caapt\_suggests\_replacement*, *caapt\_suggests\_amendment*. These take *caapt:Suggestion* as domain, *caapt:TermRoot* as range, and have no declared superproperty.
- A property which describes the relationship between a suggestion and the type of catalogue record field that the suggestion takes as part of its context: *caapt\_encountered*. This takes *caapt:Suggestion* as its domain, *skos:Collection* as its range, and has no declared superproperty. Initial label values for instances of this class are “historical context” and “contemporary context” as this is the language used in the TGD.

### 6.4 Knowledge Graph Population

The knowledge graph has been structured according to the proposed ontology and populated with the contents of the V&A TGD using a combination of manual and automated methods. The TGD was transformed by hand from its original text document form into a set of CSVs, which were then used as the input for conversion scripts.

### 6.5 Validation With Domain Experts

Validation has been carried out with V&A staff to review the conceptual model and the populated contents of the “suggestions” section of the knowledge graph. The first review was conducted using diagrams of the model and its components as the stakeholders largely hold non-technical roles and therefore an illustrative diagram was a better artifact for communication than data files or code for this audience. In addition to the model itself, an example from the TGD was shown in this structure in order to illustrate how the ontology would be used to represent a real-life example familiar to the V&A staff. An OWL file of the ontology and TTL file of the example record were also made



available so that interested parties would be able to review these materials, but the review discussions were driven by the diagrams. This review was successful, with only one change being required: in the TGD, entries often have a note saying “consult with” followed by the name of an organization external to the V&A. While that had initially been interpreted as meaning “if a question comes up regarding this term that we cannot resolve internally, this is an identified expert to consult with”, the review revealed that the meaning was in fact an internal note between the document authors and not intended as a permanent component of the TGD. As such, this relationship was removed from the modelling. As the second review was focused on contents as opposed to structure, a spreadsheet was used as the primary artifact for consideration. In order to create this spreadsheet, SPARQL was used to query the knowledge graph and return *Suggestion* labels and comments, along with relationships to other suggestions (*preferred*, *if\_not\_possible\_use*, *use\_along\_with*), grouped by *TermRoot*. The results of this review were also successful: the primary change that resulted from this review meeting was the addition of the *use\_along\_with* property between a greater number of *Suggestions*.

## 7 Conclusion

The ontology and knowledge graph have been validated by V&A staff, and the original design appears to be fit for purpose. The next step in this work will be to validate it against additional terminology guidance resources before linking it to wider LOD resources. The steps taken so far have built a solid foundation for the development of interoperable, machine-readable terminology guidelines for museums to learn from, use, or adapt to their local contexts. This project is well positioned to make a number of contributions to knowledge through the approach that is taken as well as outputs that are produced. The work that has been completed to date is already evidence of this: the gap analysis of existing LOD schemas and resulting model that makes use of existing structures where possible to enable conversations across currently loosely connected domains, along with the identification of currently unmet data representation needs, will become the first contribution to knowledge made by this project. The resulting knowledge graph has the potential to address a significant pain point for museum staff seeking to engage in critical cataloguing in both the V&A and the wider museum field.

**Acknowledgments.** This research is taking place as part of the AHRC-funded Collaborative Doctoral Partnership scheme grant AH/X004775/1.

**Disclosure of Interests.** The author has no competing interests to declare that are relevant to the content of this article.

## References

1. Chew, C.: Decolonising Description: Addressing Discriminatory Language in Scottish Public Heritage and Beyond. *Journal of Irish and Scottish Studies*. 11, 3 (2023). <https://doi.org/10.57132/jiss.213>.

2. Dalal-Clayton, A., Rutherford, A.: Against a New Orthodoxy: Decolonised “Objectivity” in the Cataloguing and Description of Artworks, <https://photoarchive.paulmellon-centre.ac.uk/groups/against-a-new-orthodoxy>, (2020).
3. Museums Association: Supporting Decolonisation in Museums, <https://www.museumsassociation.org/campaigns/decolonising-museums/supporting-decolonisation-in-museums/>, (2021).
4. Tropenmuseum: Words Matter: An Unfinished Guide to Word Choices in the Cultural Sector, <https://www.materialculture.nl/en/publications/words-matter>, (2018).
5. Chilcott, A.: Towards protocols for describing racially offensive language in UK public archives. *Archival Science*. 19, 359–376 (2019). <https://doi.org/10.1007/s10502-019-09314-y>.
6. Kennedy, L.: Finding the Words: Addressing language in archive collections as archivists, heritage professionals and volunteers, <https://bathandcolonialism.org/wp-content/uploads/2022/04/Finding-the-Words-booklet.pdf>, (2022).
7. Chew, C.: Welcome to CHTNUK, <https://culturalheritageterminology.co.uk/>, last accessed 2023/04/28.
8. Frick, R., Proffitt, M.: Reimagine Descriptive Workflows: A Community-informed Agenda for Reparative and Inclusive Descriptive Practice. OCLC Research (2022).
9. Bekiari, C., Bruseker, G., Canning, E., Doerr, M., Ore, C.-E., Stead, S., Velios, A.: CIDOC Conceptual Reference Model, Version 7.1.2, <https://cidoc-crm.org/Version/version-7.1.2>, (2022).
10. Khan, F., et al.: When Linguistics Meets Web Technologies. Recent advances in Modelling Linguistic Linked Open Data. *Semantic Web*. 0, 1–62 (2021).
11. Hardesty, J., Nolan, A.: Mitigating Bias in Metadata: A Use Case Using Homosaurus Linked Data. *Information Technology and Libraries*. 40, (2021). <https://doi.org/10.6017/ital.v40i3.13053>.
12. Nesterov, A., Hollink, L., van Erp, M., van Ossenbruggen, J.: A Knowledge Graph of Contentious Terminology for Inclusive Representation of Cultural Heritage. In: Pesquita, C., Jimenez-Ruiz, E., McCusker, J., Faria, D., Dragoni, M., Dimou, A., Troncy, R., and Hertling, S. (eds.) *The Semantic Web*. pp. 502–519. Springer Nature Switzerland, Cham (2023). [https://doi.org/10.1007/978-3-031-33455-9\\_30](https://doi.org/10.1007/978-3-031-33455-9_30).
13. Khan, F., Salgado, A.: Modelling Lexicographic Resources using CIDOC-CRM, FRBRoo and Ontolex-Lemon. Presented at the SWODCH: Semantic Web and Ontology Design for Cultural Heritage (2021).
14. Canning, E., Brown, S., Roger, S., Martin, K.: The Power to Structure : Making Meaning from Metadata Through Ontologies. *KULA: Knowledge Creation, Dissemination, and Preservation Studies*. 6, 1–15 (2022). <https://doi.org/10.18357/kula.169>.
15. Cummings, J., Stacey, D.: Lean Ontology Development. In: *Proceedings of the 10th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*. pp. 367–374 (2018). <https://doi.org/10.5220/0006963003670374>.
16. Watson, B.M.: Critiquing the Machine: The Critical Cataloging Database. *TCB: Technical Services in Religion & Theology*. 31, 1–17 (2023). <https://doi.org/10.31046/tcb.v31i1.3216>.