

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

Erin Canning¹[0000-0003-3383-272X]

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

Abstract. 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 problematics in their 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, already familiar to many museums through the datafication of collections information. The use of linked open data will allow for the creation of a knowledge graph that can link together machine-readable versions of terminology guidance from multiple sources and then further connect them to additional resources such as commonly used vocabularies like the Getty Art & Architecture Thesaurus, which can be found in the controlled vocabularies of many museums. This paper describes the process of developing an ontology for the representation of cultural heritage terminology guidelines documents, using the Victoria & Albert Museum’s Terminology Guidelines document as a case study. The resulting knowledge graph contributes to the understanding of how museums are thinking about what makes a term problematic and what possible reparative actions are given different contexts, and supports museum workers engaged in this critical cataloguing work by addressing the need to develop sector-wide guidance and support the sharing of resources across institutions.

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

1 Introduction and Motivation

1.1 Introduction

The Victoria & Albert Museum (V&A) first opened its doors to the public as the South Kensington Museum in 1857. The catalyst for the creation of the institution can be traced to the 1851 Great Exhibition of the Industry of All Nations: the organization of the Great Exhibition was led by Prince Albert and Henry Cole, who would become the first director of the museum, profits from admission tickets were used to purchase the

museum site, and the original core collection of museum objects were purchased from the Great Exhibition itself [1]. The museum was renamed to the Victoria & Albert Museum in 1899, and today it contains close to 1.7 million works of art and design objects, organized across five collecting departments: Art, Architecture, Photography & Design; Decorative Art and Sculpture; Performance, Furniture, Textiles and Fashion; Asia Department; and Young V&A, formerly called the Museum of Childhood. The V&A Museum has an extraordinarily diverse collection, containing objects collected over its 170-year history from all reaches of the British colonial empire.

Just as diverse as the objects are the cataloguing records that accompany them: the V&A Museum catalogued their acquisitions from the origins of the museum and these records have been translated from analogue forms of accession books and paper records through computerization to the technical systems that are used to manage them today, known as collections management systems. 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, from the information fields available in the database, which represent what kinds of information have been considered important to record, to 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 Museum, and the museum field as a whole is wrestling with how to best address the problematics in their data.

1.2 Motivation

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 problematics 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 [2–4]. There is an opportunity to address this through the use of linked open data (LOD), already familiar to many museums through the datafication of collections information. Curators and cataloguing staff 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 allow for the creation of a knowledge graph that can link together machine-readable versions of terminology guidance from multiple sources and then further connect them to additional LOD resources such as commonly used vocabularies like the Getty Art & Architecture Thesaurus (AAT), which can be found in the controlled vocabularies of many museums. The development of such an ontology will contribute to the understanding of how museums are thinking about what makes a term problematic and what possible reparative actions are given different contexts, as well as supporting museum workers engaged in this critical cataloguing work by addressing the need to develop sector-wide guidance and support the sharing of resources across institutions.

2 Related Work

2.1 Addressing Problematic Terminology in Museum Catalogue Data

The Museums Association “Decolonization Guidance Working Group” produced the “Supporting decolonization in museums” guidelines in 2021, two pages of which are dedicated to collections cataloguing [5]. Resources aimed at the museum community have also been developed by individual people and institutions and while these come from places of localized work, the findings are nonetheless applicable for others. “Words Matter: An Unfinished Guide to Word Choices in the Cultural Sector”, a document written by the National Museums for World Cultures, includes a glossary of problematic terminology and suggestions for what other museums could do when encountering the terms in their own catalogue data [6]. In the UK, Alicia Chilcott published a set of recommendations for conducting critical cataloguing work and the “Bath and Colonialism Archive Project” produced a guide for addressing problematic terminology in archival catalogue data following, and building on, Chilcott’s recommendations [7, 8]. Carissa Chew developed the Cultural Heritage Terminology Network Glossary, a collaborative resource intended to “promote cross-institutional collaboration on inclusive description issues” [2]. In all of these cases, although the resources started locally, they respond to a field-wide need and seek to contribute to a collective effort.

Individual institutions are also 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 replaced, marked as historical, marked as a historical with a new term appearing alongside, or kept but contextualized. There is no single rule that would be appropriate in all cases—even for in a given institution—and it is essential to take the context and purpose of the work into careful consideration [3, 9, 10]. 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

Identifying Existing Ontologies. 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 of ontologies are the CIDOC Conceptual Reference Model (CIDOC CRM) for its coverage of cultural heritage data, the Ontology Lexicon (OntoLex) suite of ontologies for its coverage of linguistic data, the Simple Knowledge Organization Schema (SKOS) for its coverage of thesauri and used in critical cataloguing LOD vocabulary projects such as Homosaurus, and the Cultural Contexts Concept Scheme for Contentious Terminology (CULCO).¹ These ontologies are introduced here, and evaluated as part of the gap analysis work described below in section 6.2.

- CIDOC CRM is an ontology for the representation of cultural heritage data [11]. It is a stable standard widely used in the cultural heritage community and has been an ISO standard since 2006, most recently updated in 2023.
- OntoLex is an ontology for the representation of lexica and dictionaries. It is a stable standard, and the most widely used model for publishing these linguistic resources as linked open data [12].
- SKOS is 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.
- CULCO is an ontology developed in 2022 as part of the “Culturally Aware AI” project to describe the glossary section of the “Words Matter” publication [13].

Connections Between Existing Ontologies. These four identified ontologies do not exist in isolation. Points of connection link these structures 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 a key ontology for bringing these schemas together. Further connections can be found in [14], in which the authors propose an alignment between OntoLex and CIDOC CRM. While these connections are usually encoded, others are more informally proposed in scope notes or wider documentation.

3 Problem Statement and Contributions

Cultural heritage institutions are struggling with taking initial steps in locating and addressing problematic terminology in their catalogues. The lack of domain-wide guidance and scarcity of shared, or shareable, resources regarding problematic terminology and possible reparative actions are significant barriers for museum professionals seeking to engage in this work. A linked open data solution that makes use of an ontology for the representation of terminology guidance documents – specific terms and forms, the factors through which the terms are defined by museums, and suggestions on reparative actions – 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.

¹ CIDOC CRM: <https://cidoc-crm.org/>; OntoLex: <https://www.w3.org/2016/05/ontolex/>; SKOS: <https://www.w3.org/2009/08/skos-reference/skos.html>; Homosaurus: <https://homosaurus.org/>; CULCO: <https://cultural-ai.github.io/wordsmatter/>

The novel outputs that will be produced are: an ontology for representing problematic terminology and suggested actions, taxonomies of types of problematic terminology and types of suggested actions, and a populated and linked knowledge graph. This work will also lead to the ability to compare between institutions what is being looked for (categories of offensive language), how it is being identified (terms and forms), and what is being recommended when instances are encountered, taking into account various contextual elements (suggestions and categories of suggestions).

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 & Albert Museum’s 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 intended 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.

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 the “Words Matter” publication, and the Cultural Heritage Terminology Network Glossary (CHTNG) [2, 6]. These are frequently cited resources: “Words Matter” is regularly looked to as a guiding document in this field, and the CHTNG 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 Getty AAT and Homosaurus – a community-developed LOD vocabulary that is already working in the space of critical cataloguing – to assert relationships to terms published by those sources. It is anticipated that terms that appear in the knowledge graph will also be part of these vocabularies, which offers a route to wider data interoperability. This will also work to expose how problematic terminology may be introduced into collections data through the use of externally-provided vocabularies, as well as how community-developed vocabularies may offer alternative solutions.

5 Evaluation Plan

Evaluation will be done in three parts: through review with V&A Museum 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 Museum staff. These are staff who are responsible for creating and

maintaining the museum's Terminology Guidance Document, and therefore are best positioned to evaluate whether or not 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 for this section is therefore critical for the creation of the knowledge graph.

Integrating the "Words Matter" and CHTNG glossaries 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 Museum staff to ensure that the base model being compared is a stable foundation for this stage of the work.

Lastly, validation through competency questions will be performed once the two additional terminology guidelines have been integrated, and links to the two additional vocabularies sources have been made. These questions will be developed in collaboration with the V&A Museum 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 ongoing critical cataloguing work.

6 Results

6.1 Conceptual Model of the Victoria & Albert Museum Terminology Guidance Document

The V&A Museum has formally held cross-departmental meetings to discuss terminology concerns since July 2016. These meetings provide a space for staff members to discuss problematic terminology they have encountered in museum catalogue records, with the goal of coming to recommendations about how to proceed in general and specific cases. Decisions are documented and inform the V&A Museum 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 of information: the potentially problematic term to be considered, a description of the term, and suggestions for cataloguing staff to take when they encounter the term in cataloguing data. Within each of these sections is a wealth of information:

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
3. Suggestions - a list of suggestions that could be considered based on specifics of the term encounter, such as: what place is being described in the record, what group of

people is being described of the record, the original intended use of the term in the record, and the field in the record where the term is found

A suggestion may recommend that a different term be used to replace or augment the potentially problematic term under consideration. Therefore, the TGD contains information about not only potentially problematic terms, but also potentially preferable ones. As such, a conceptual model of the document needs to accommodate the representation of the classification of terms as potentially problematic or preferable.

Lastly, there is also information about the TGD itself to be considered: 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.

The conceptual modelling stage of work separates each of these elements into separate entities and defines the relationships between them. Despite having discrete information considerations, each section is connected together, most notably by shared reference to the people and places that are described when terms are used.

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 contents of the terminology guideline document “Words Matter”: this document includes a glossary of terms accompanied by a description of their history, use, and possible sensitivities, as well as suggestions of actions for cataloguers to take when they encounter the term in a museum record [6]. 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 encountered by cataloguers – 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 (described in [15]). As such, it was found to be a strong foundation for future ontology development work, but not sufficient on its own.

OntoLex, on the other hand, introduces a level of complexity in modelling lexical information beyond what is required for the representation of the TGD: exploring using OntoLex for representing the terminology form and description information would require the use of classes such as *LexicalConcept* 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—the objects about which the records under consideration have been written—that

6.4 Knowledge Graph Population

The knowledge graph has been populated with the contents of the V&A Museum 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. This intermediary step was invaluable for ensuring transparency, especially when communicating with non-technical stakeholders.

6.5 Validation With Domain Experts

Validation has been carried out with V&A Museum staff to review the conceptual model and translation of the “suggestions” section. The first 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 Museum. 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. The results of the second review were similarly 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 Museum 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 at both the V&A Museum and in the wider cultural heritage community.

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. Barringer, T.: The South Kensington Museum and the colonial project. In: Barringer, T. and Flynn, T. (eds.) *Colonialism and the Object*. pp. 11–27. Routledge, London (1997).
2. 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>.
3. 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>.
4. Museums Association: Our statement on Decolonisation, <https://www.museumsassociation.org/campaigns/decolonising-museums/our-statement-on-decolonisation/>, last accessed 2023/04/28.
5. Museums Association: Supporting Decolonisation in Museums, <https://www.museumsassociation.org/campaigns/decolonising-museums/supporting-decolonisation-in-museums/>, last accessed 2023/04/28.
6. Tropenmuseum: Words Matter: An Unfinished Guide to Word Choices in the Cultural Sector, <http://www.tropenmuseum.nl/en/about-tropenmuseum/words-matter-publication>, (2018).
7. 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>.
8. 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).
9. Frick, R., Proffitt, M.: Reimagine Descriptive Workflows: A Community-informed Agenda for Reparative and Inclusive Descriptive Practice. OCLC Research (2022).
10. Rutherford, A.: Provisional semantics, projects and positionality. *ARLIS Cataloguing and Classification Ethics Series 2022*. (2022).
11. Bekiari, C., Bruseker, G., Doerr, M., Ore, C.-E., Stead, S., Velios, A.: CIDOC Conceptual Reference Model, V. 7.1.1 (2022). <https://doi.org/10.26225/FDZH-X261>.
12. Khan, F., et al.: When Linguistics Meets Web Technologies. *Recent advances in Modelling Linguistic Linked Open Data*. *Semantic Web*. 0, 1–62 (2021).
13. 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.
14. 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).
15. 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>.