

# Evaluating end-to-end entity linking on domain-specific knowledge bases: Learning about ancient technologies from museum collections

**Sebastian Cadavid-Sanchez**  
Sciences Po

**Khalil Kacem**  
Sciences Po

**Rafael Aparecido Martins Frade**  
Sciences Po

**Johannes Boehm**  
Sciences Po

**Thomas Chaney**  
USC & Sciences Po

**Danial Lashkari**  
Boston College

**Daniel Simig\***

## Abstract

To study social, economic, and historical questions, researchers in the social sciences and humanities have started to use increasingly large unstructured textual datasets. While recent advances in NLP provide many tools to efficiently process such data, most existing approaches rely on generic solutions whose performance and suitability for domain-specific tasks is not well understood. This work presents an attempt to bridge this domain gap by exploring the use of modern Entity Linking approaches for the enrichment of museum collection data. We collect a dataset comprising of more than 1700 texts annotated with 7,510 mention-entity pairs, evaluate some off-the-shelf solutions in detail using this dataset and finally fine-tune a recent end-to-end EL model on this data. We show that our fine-tuned model significantly outperforms other approaches currently available in this domain and present a proof-of-concept use case of this model. We release our dataset and our best model.

## 1 Introduction

Digital technologies are enabling social scientists to easily access large quantities of textual information. These large textual datasets are also increasingly used by empirical researchers as a substitute for traditional data in numerical form (Gentzkow et al., 2019). The main challenge in doing so lies in the way that textual data is coded into numerical form. Traditionally researchers relied on experts performing the coding procedure manually, often at great expense.<sup>1</sup> While in the past years several Entity Linking (EL) systems have been developed

(Kolitsas et al., 2018; Gillick et al., 2019; Wu et al., 2020; Li et al., 2020) that could possibly provide much better solutions for social scientists, these models are often developed and evaluated on generic knowledge bases (e.g. Wikidata) and datasets that are not representative of the challenges social scientists face.

This work aims to enable researchers in this field to better leverage the advances in NLP. We first present a dataset we collected for the purpose of finetuning and evaluating end-to-end entity linking systems on descriptions of various museum artifacts, then use this dataset to explore the performance of different entity linking solutions at tagging and linking entities to a domain-specific knowledge base from the digital humanities. Our Knowledge Base is the [Getty Arts and Architecture Thesaurus](#), a controlled vocabulary describing terms associated with cultural heritage objects (predominantly materials, processes, objects, and attributes). We find that our preferred model is able to greatly outperform naive baselines currently available to researchers in the Social Sciences field.

We believe that our work contributes both to the NLP community by evaluating the generalizability of entity linking methods to different knowledge bases, as well as to the social sciences by evaluating the effectiveness of modern end-to-end entity linking methods for nuanced classification problems. We also contribute to the digital humanities and information sciences, where researchers and practitioners frequently work with controlled vocabularies and face classification problems, but are reluctant to use automated techniques like machine learning (Coeckelbergs and Van Hooland, 2021) and

\*Correspondence to: simigd@gmail.com

<sup>1</sup>As an example, the Human Relations Area Files (Murdock et al., 2008) is a curated database of text paragraphs from the cultural anthropological literature, where experts have classified

each paragraph by culture and subject. The project was started in the 1940's, and is still ongoing.

often prefer rule-based methods or crowd-sourcing (Hall and Walsh, 2021; Ridge, 2014). We offer an off-the-shelf solution that can be used and further developed by practitioners.

## 2 The ARTIFACT Dataset

### 2.1 Sourcing Museum Labels

To create this dataset we use textual descriptions of objects from six online museum collections<sup>2</sup>. Each object in these collections is annotated with 4 different text fields: title, detailed description, and free-form metadata about materials and techniques. Table 1 provides examples.

In some cases, different types of text fields may be repetitive in terms of entity content and contain texts of different complexity. In order to efficiently use our resources and avoid redundancy, we sample from different types of text fields independently.

Type	Example
Title	"Fragment of Carved Wood Panel with Two Birds"
Description	"Fragment. With part of figure of ghulám, sleeve. Made of black, grey-brown, blue painted plaste"
Material	"Gilt-copper alloy, rock crystal"
Technique	"Plain weave, tapestry, dyed"

Table 1: Examples for each object description type in our digital museum data

To evaluate the performance of different Entity Linking techniques, we take a small sample from this dataset and manually annotate it with EL labels. Since some materials and objects appear much more frequently than others, we increase the variability of labels in our sample by implementing a stratified sampling process based on the registered date for each object. We provide more details on the sampling process in Appendix A.1.

<sup>2</sup>British Museum, Metropolitan Museum of Art, Williams College Museum, Cleveland Museum of Art, Penn Museum and the Smithsonian Museums. The data in this paper is part of a broader project on ancient technologies, using tens of millions of objects from dozens of museums

### 2.2 Using a domain-specific knowledge base

Unlike most Entity Linking applications that use Wikipedia as their knowledge base, our dataset links entities to the *Getty Arts and Architecture Thesaurus* (AAT),<sup>3</sup> an expert-curated taxonomy of more than 50,000 concepts related to arts, architecture and visual cultural heritage. All concepts are structured hierarchically in eight facets (subtrees) (J. Paul Getty Trust, 2022a). Each entry has a definition and belongs to a rank in the hierarchy structure. We focus on four facets of the taxonomy, listed in Table 2.

Hierarchy	Examples
Objects	Knives, vessels, amphorae, neck, earthenware, handle.
Materials	Ceramic, gold, clay, wood, calcite, diamond, glass.
Styles and Periods	Periods Peruvian, ancient chinese, greek, gothic, inca.
Activities	Machine sewing, sealing, gilding, hammering.

Table 2: Examples for analyzed hierarchies of the taxonomy

AAT is different from Wikipedia in many ways:

- AAT focuses on a small subset of domains covered by Wikipedia but provides a much more granular coverage of entities in those domains. For example, a search in Wikipedia for the word "hammer" returns a few different results, while AAT contains more than one hundred different types of hammers. The main differences between types lie in distinctive characteristics, such as specific use, material, period of manufacture, etc.
- Unlike in the case of Wikipedia that has a hyperlink graph, in our setting one cannot leverage pre-computed mention tables to aid candidate generation.
- Descriptions in AAT are much shorter than in Wikipedia—their purpose is merely to provide a

<sup>3</sup><https://www.getty.edu/research/tools/vocabularies/aat/about.html>

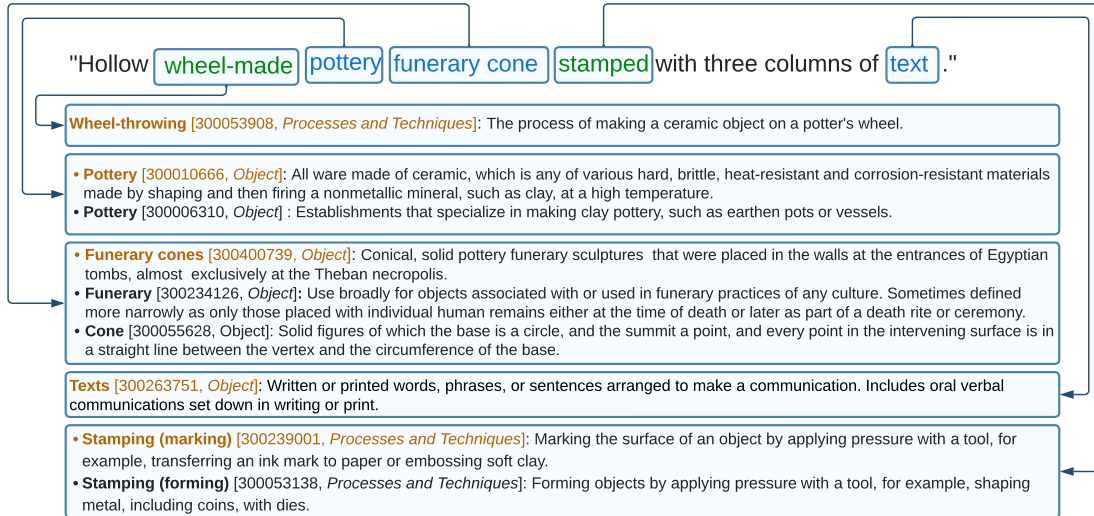


Figure 1: Example of end-to-end entity linkage of a museum item description. Each box contains recognized entities (Mention Detection). The green text represents the Process and Techniques facet, and the blue text represents the Object facet. The light-blue boxes include the most relevant candidates for entity linking (Candidate Generation), and the gold text indicates the correct candidate (Entity Disambiguation). Each integer identifier represents the corresponding AAT ID for each candidate.

definition of the entity in question. On average, AAT descriptions are two sentences long.

- In our dataset we focus solely on non-named entities and ignore any named entities - in contrast with a large body of literature focused on named entities in Wikipedia.

To the best of our knowledge, all prior approaches used by scholars in the digital humanities to link text to AAT entities involve manual annotation, in some cases supported by string and keyword matching (J. Paul Getty Trust, 2022b; Smithsonian Institution, 2022).

### 2.3 Annotation Process

To produce manually-annotated data, sampled strings were distributed evenly among 4 annotators. Instruction lessons were conducted to show the annotators gold examples and orient them to use an in-house tool to produce and store annotations.

Throughout the annotation process, a quality control exercise was repeatedly carried out. In this the 4 annotators received an identical, small batch of data, annotated them independently, then discussed their results together and recorded the consolidated labels. This process ensured that

	Metric (%)	Upper bound	Median	Lower bound
MD	Precision	94.8	92.4	88.9
	Recall	90.6	80.9	71.4
	F1	92.2	86.5	79.2
ED	Precision	88	81	71
End-to-end	Precision	88.3	85.6	82.9
	Recall	85.3	75.8	67.9
	F1	86.8	80.4	74.7

Table 3: Performance statistics on the jointly labeled data. We consider the commonly agreed label as the gold label and measure strong-matching performance of individual raters' against that.

labels from different raters are consistent, but also resulted in a smaller dataset (552 annotated mention-entity pairs) that we use as a high quality test set for our evaluations. Table 3 shows performance statistics of individual raters when compared to the commonly agreed-upon ground truth. Notably, we found that there exists some amount of unavoidable ambiguity in the Entity Linking task due to the limited information available in descriptions - even our best performing rater could only produce and end-to-end F1 score of 86.8 when compared to the consensus labels.

The final, single-annotated set comprises of 1,728 strings annotated with a total of 7,510 entity

labels. This annotation process (including discussion sessions) took approximately 300 hours of human labor. Given the difficulty of the Entity Linking task, it was initially costly to standardize criteria between annotators for disambiguating between different entities in different contexts.

We use the single-annotated dataset to finetune models for our domain. We randomly split this dataset into training (90%) and validation (10%) sets. Main statistics of these data sets are presented in Table 7 and additional details are provided in Appendix A.2. Our dataset called ARTIFACT<sup>4</sup> is available at: <https://github.com/LongRunGrowth/ARTIFACT-Dataset>.

## 2.4 Challenges of Entity Linking in digital museum data

We highlight several Entity Linking challenges presented by our dataset, using the example in Figure 1. We believe these properties make ARTIFACT a suitable dataset to finetune and evaluate generic EL systems on this specific domain.

**Context-dependent disambiguation** In 43.5% of the cases, matching semantic context from the text and the entity definition is crucial to make the correct selection, as multiple entities with similar names could be a possible match based on the surface form.<sup>5</sup> In Figure 1, the mention "pottery" has two AAT candidates with identical names; however, their descriptions vary substantially. Using the context suggests the text does not refer to a workshop where pottery is made, but to a pottery object. A similar case applies to the word "stamped" in the example.

**Surface form diversity.** 9.5% of the labels in our dataset have a significantly different entity name and mention text. This is a failure case for the currently prevalent fuzzy-matching based methods and currently requires an annotator to do a detailed search or know the entity in advance. This is

<sup>4</sup>Automated Resource for Technology and Information Finding in Artifacts CollecTions.

<sup>5</sup>To calculate this metric, we sampled over 30% of errors generated by our string-matching baseline (described in section 3) and classified them into context-dependent and non-context-dependent errors. Using this metric, we approximate the proportion of context-dependent disambiguation cases over our full annotations set.

illustrated in Figure 1 for the "wheel-made" mention: usually, this term refers to pottery made using a wheel. However, "wheel-made" does not match exactly with any entity in our knowledge base, but a similar expression with the same meaning exists, "wheel-throwing". Using a simpler linking algorithm, such as string matching, this type of mention could be mapped to the incorrect entity "wheel making", which corresponds to "the art of making wheels".

**Overlapping Mentions** This is exemplified in Figure 1 by the mention "funerary cone". Although there is an entity that matches exactly the full mention, an annotator may tag "funerary" and "cone" and link them to separate entities.

**Incomplete Taxonomy** We note that there were cases where even though a referenced entity belonged to one of the analyzed facets, it was not currently included in AAT and thus could not have been included in our dataset. The incompleteness of taxonomies for high-cardinality classification tasks is a known problem (Simig et al., 2022), and addressing it is beyond the scope of this work.

## 3 Baseline Systems

We provide and evaluate a set of baseline EL solutions along with the dataset, covering a wide range of system complexities. We first test model-free approaches currently available, as well as off-the-shelf machine learning based systems. Finally, we adapt a more recent model originally trained on Wikipedia (ELQ, Li et al. 2020) to our domain.

### 3.1 Parameter-free Baselines

**Naive String Matching** We build a pretrained SpaCy (Honnibal et al., 2019) pipeline that provides an end-to-end framework by integrating an i) "Entity-Recognizer"<sup>6</sup> model to predict non-overlapping mention boundaries using fuzzy matching to search spans from entities in the knowledge base and an ii) "Entity-Ruler"<sup>7</sup> model to perform Candidate Generation and Entity Disambiguation tasks. The model performs random candidate disambiguation when there are similar matches between candidates.

<sup>6</sup><https://spacy.io/api/entityrecognizer>

<sup>7</sup><https://spacy.io/api/entityruler>

Using the SpaCy-pipelines platform allows users wide accessibility and fast usage, and to the best of our knowledge, is the most accessible tool for entity-linking tasks on museum data before our work.

**Memorization** We establish a rule-based entity disambiguation baseline in which we link a mention in the test set to the entity most commonly associated with this mention in the train set. For example, the string "buff" could be tagged to two different entities (a tool or a pottery style). This ruled-based baseline would tag all occurrences of "buff" with the most common entity in the train set. We expect systems that truly utilize context (as opposed to just overfitting on the seen mention-entity pairs) to outperform this baseline significantly.

### 3.2 Model-Based Disambiguation

**SBERT** (Reimers and Gurevych, 2020) computes text embeddings that can be used to identify similar sentences. As a simple baseline, we use our AAT Knowledge Base to create embeddings for each entity description and similarly embed the text in which the mention appears. We choose the entity with the highest cosine similarity.

**GENRE** (De Cao et al., 2021) is a sequence-to-sequence model that uses BART (Lewis et al., 2019) to link tagged mentions to entities from Wikipedia using constrained beam search to generate valid identifiers.

In order to link the output of GENRE to AAT entities, we build a database that maps Wikipedia IDs and labels with their corresponding AAT when it exists. When the Wikipedia label is not specific enough and corresponds to different AATs (as in the example provided in subsection 2.2), a candidate is randomly selected. We expect this to be a weak baseline that highlights the limitations of systems that rely on Wikipedia as their knowledge base.

### 3.3 An End-to-End Solution

**ELQ** (Li et al., 2020) is an entity linking system that performs both mention detection and disambiguation, originally developed to help question-answering models. ELQ detects mentions and computes mention representations in a single pass of inference, retrieves entity candidates using a K-Nearest Neighbor (KNN) search and chooses the best candidate based on embedding similarity. We use

this model as our strongest baseline as we found it to provide the best trade-off between low system complexity and good performance on standard standard EL benchmarks.

We concatenate names and descriptions of AAT entities to form the text to be encoded by ELQ’s bi-encoder. We first evaluate an off-the-shelf version of the ELQ model provided by Li, Min, Iyer, Mehdad, and Yih (2020), then also finetune this on our training set in order to adopt it to our domain specific texts and entities. We run all training exercises for 500 epochs using a batch size of 256 on a single NVIDIA A100 GPU. We use AdamW optimizer with learning rate 1e-5 and linear decay.

## 4 Evaluation of Baseline Systems

### 4.1 Evaluation Protocol

We separate evaluation into three sections: Mention Detection, Entity Disambiguation, and End-to-End performance. To evaluate mention detection, we use strong matching: an annotation is correct if its boundaries match the ground-truth mention boundaries in the text exactly. When evaluating entity disambiguation, an annotation is correct if the linked entity ID for a given mention precisely matches the AAT ID in the ground-truth annotation. An End-to-End annotation is considered correct only if it matches both the mention boundaries and AAT ID of the ground-truth annotation exactly.

In order to compare Entity Disambiguation performance among a wide range of disambiguation solutions (Table 5), we select as input the correctly tagged mentions from the ELQ model. For this experiment we further divide tagged items into (i) *Seen* tags that are present in annotations of both the training set and the test set (88.94 %) and (ii) *Unseen* Tags that did not appear in the training set (11.06%). The latter set - albeit small - allows us to evaluate the generalization capacity of our models when considering unobserved tags, which could appear considerably in our application case.

### 4.2 Results

The results of evaluation for Mention Detection, Entity Disambiguation and End-to-End Entity Linking on our held out, high quality test set are presented in Tables 4, 5 and 6 respectively.

Annotation-system	Model version	Metric (%)		
		Precision	Recall	F1
String Matching	Off-the-shelf	57.6	42.9	49.1
ELQ	Off-the-shelf	62.9	44.2	51.9
ELQ	Fine-tuned	<b>85.4</b>	<b>81</b>	<b>83.1</b>

Table 4: Performance for annotation systems able to perform Mention Detection.

Annotation-system	Model version	% Precision		
		All tags	Unseen	Seen
String-Matching	Off-the-shelf	45.5	49.8	44.5
Memorization baseline	-	78.6	0	87.4
SBERT	Off-the-shelf	64.6	<b>74.6</b>	62.7
GENRE	Off-the-shelf	30.5	43.1	28.7
ELQ	Off-the-shelf	39.9	43.4	8.2
ELQ	Fine-tuned	<b>88.7</b>	68.4	<b>90.2</b>

Table 5: Performance for annotation systems able to perform Entity Linking.

Annotation-system	Model version	Metric (%)		
		Precision	Recall	F1
String Matching	Off-the-shelf	59.6	39.5	47.5
ELQ	Off-the-shelf	57.7	42.2	48.7
ELQ	Fine-tuned	<b>78.5</b>	<b>74.6</b>	<b>76.9</b>

Table 6: Performance for annotation systems able to perform entity linking End-to-End.

Even a state-of-the-art model such ELQ performs only on par with simple baselines when not adapted to the domain of our texts and the custom knowledge base, AAT. However, fine-tuning ELQ even on our relatively small training dataset results in a step change in the model’s performance. Notably, Entity Disambiguation performance on the *Unseen* set also improves, demonstrating that the ELQ model learns to generalize from the training set we provide. Interestingly, the off-the-shelf SBERT solution for disambiguation outperforms finetuned ELQ when linking unseen AATs, indicating that there is still room for improvement for future end-to-end models. Despite that, the fine-tuned ELQ model reaches the lower-bound of human performance metrics for Mention Detection and End-to-end tasks.

## 5 An Example Use Case

Figure 2 shows, as a proof-of-concept example, that entity linking can a powerful tool to study the invention and diffusion of technologies over time and space. We apply our End-to-End entity

linking system to 2.7 million objects from 26 English-language museum collections. We define an object as ‘technologically significant’ (large radius) if it belongs to an important technology class (a class with many objects) and if it is a pioneer within its technology class (early discoveries).

The map allows to visualize the process of technological diffusion over more than ten millennia. Neolithic sites with significant technologies (large blue circles) appear along the Nile and Indus rivers, Southern Turkey and the fertile crescent, as well as in Southwest France. Other less significant Neolithic sites are scattered over Europe, the Middle East, and South Asia (small blue circles). Over the Bronze and Iron Ages (purple), significant technologies are concentrated along the Nile, the Indus, and the fertile crescent. Classical Antiquity (pink circles) sees a shift of significant technologies to Greece, the Near East, and the Southern Italian peninsula. In the Middle Ages and the Renaissance (red and yellow circles), significant technologies are concentrated in Western Europe and South Asia.

Crucially, our ability to link individual objects to AAT entities is key to identifying ‘significant’ technologies. For instance, focusing on the Renaissance (yellow circles), while objects appear throughout Europe, the most significant discoveries (large yellow circles) are concentrated in Northern Italy, and Northern and Central Europe, as well as in important urban centers such as Paris, London, and Madrid.

## 6 Related Work

### 6.1 Entity Linking for Sciences

Entity Linking has proved to be a powerful tool in various fields of science. D’Souza, Hoppe, Brack, Jaradeh, Auer, and Ewerth (2020) present the STEM-ECR dataset with annotations to identify and classify scientific entities from abstracts of publications drawn from 10 STEM (Science, Technology, Engineering, and Medicine) disciplines. Applications of EL in biomedicine (Mohan and Li, 2019; Vashishth et al., 2020; Liu et al., 2021) and social network analysis (Derczynski et al., 2015; Adjali et al., 2020) have also been explored.

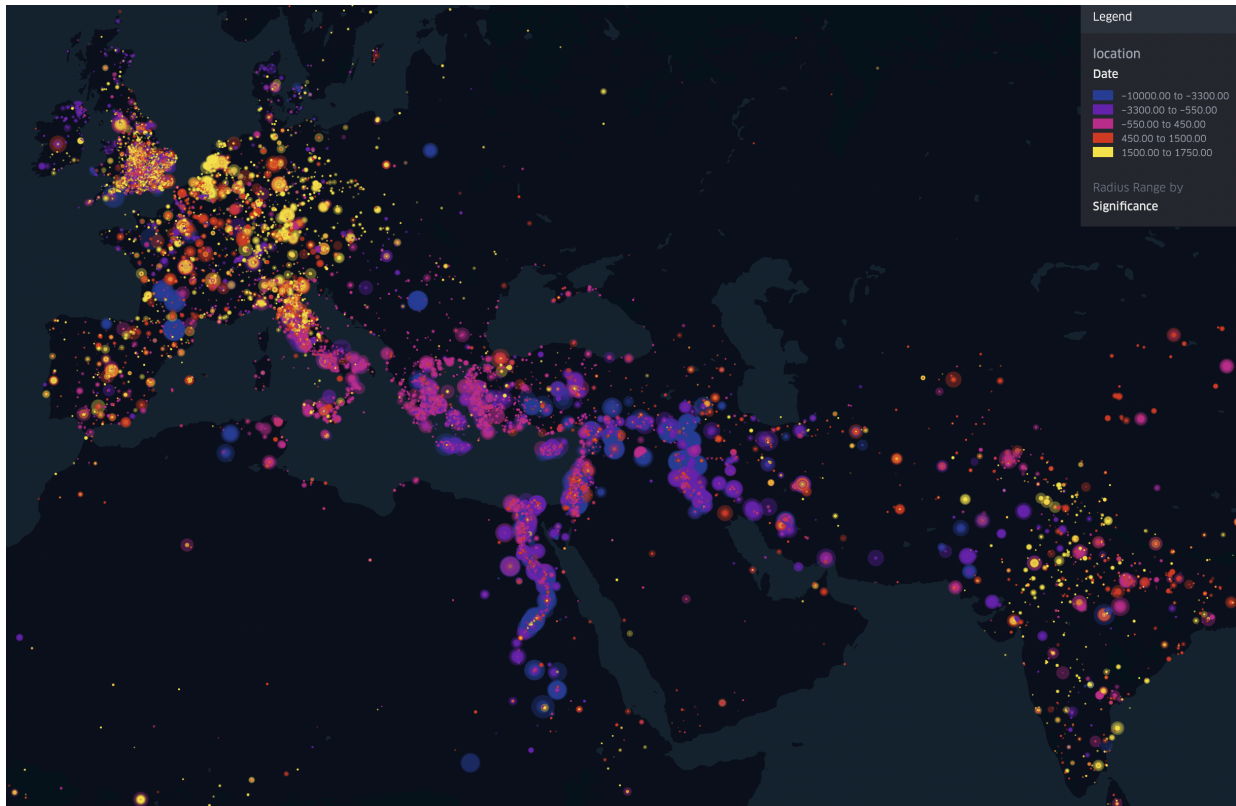


Figure 2: This figure maps 2.7 million objects from 26 English-language museum collection, where we classify objects into AAT technology categories using our End-to-End Entity Linking system (ELQ fine-tuned). A circle’s color corresponds to the date when objects are made (blue: Neolithic; purple: Bronze and Iron Ages; pink: Classical Antiquity; red: Middle Ages; yellow: Renaissance). A more opaque circle corresponds to more objects in the same location. The radius of each circle corresponds to a measure of the ‘technological significance’ of an object, constructed using our End-to-End classification. For object  $o$  that belongs to AAT technology category  $c$ :  $significance_o^c = \ln(Number_c) / Rank_o^c$ , where  $Number_c$  is the total number of objects in technology category  $c$ , and  $Rank_o^c$  is the temporal rank of object  $o$  among all objects in technology category  $c$  ( $Rank_o^c = 1$  if object  $o$  is the earliest in technology category  $c$ ,  $Rank_o^c = 2$  if it is the second earliest, etc).

## 6.2 NLP for Social Sciences

Weichselbraun, Waldvogel, Fraefel, van Schie, and Kuntschik (2022) introduce the CarrerCoach Dataset, which contains annotations across 75 documents related to education and labor market entities; and Tsereteli, Kartal, Ponzetto, Zielinski, Eckert, and Mayr (2022) present the SV-Ident dataset with annotations to perform survey variable identification in social science publications. Gillick, Kulkarni, Lansing, Presta, Baldrige, Ie, and Garcia-Olano (2019) introduces an Entity Linking dataset for news article analysis. To the best of our knowledge, our dataset is the first to provide a manually annotated database intended for quantitative social sciences research using museum collections data.

We see two broad fields of applications of our work. One is in the quantitative social sciences, where textual information is increasingly being used to construct numerical data that is being processed using statistical and econometric techniques (Gentzkow et al., 2019, Wilkerson and Casas, 2017, and Jensen et al., 2022 give overviews of such approaches in economics, political science, and sociology, respectively). So far, much of the work in this domain uses unsupervised learning methods, in particular topic modelling, to reduce the dimensionality of the data (e.g. Hansen et al., 2018, Mueller and Rauh, 2018). When the categories that observations should be classified into are known ex ante, these approaches become difficult to use.

That is the case in particular in economics, where there is often a consensus on the use of certain industry classifications or patent codes. Related to our application, the use of machine learning techniques is particularly promising in the field of economic history, where much of the original data is in textual form (see, e.g. Feigenbaum, 2016, Combes et al., 2022, or Shen et al., 2020 for examples).

The second field of applications are the digital humanities, where researchers and practitioners are often confronted with large quantities of unstructured text data. Entity linking provides a way of classifying and thereby linking records from these corpora, enhancing the usefulness and discoverability of these datasets (Knoblock et al., 2017). Our work may be of particular interest to those working on cultural heritage data, where the AAT is a commonly used controlled vocabulary (see, e.g. Van Hooland et al., 2013, or the Linked Art project, described in Delmas-Glass and Sanderson, 2020). Previous approaches in the digital humanities community either use manual entity linking (e.g. Van Hooland et al., 2013) or use approaches based on term frequencies (e.g. Mendes et al., 2011). Munnelly and Lawless (2018) describe the process of creating a domain-specific knowledge of notable Irish people and constructing an entity linker via DBpedia. Van Hooland et al. (2015) evaluate different web-based NEL services for the linking of museum collection metadata to Wikidata, but find recall rates of below 50% using these services. Brando et al. (2016) present a domain-specific graph-based entity linking strategy for historical figures. Most closely related to our approach, Linhares Pontes et al. (2022) train an end-to-end entity linker based on MUSE (Conneau et al., 2017) for linking entities in historical press articles to Wikipedia entries.

## 7 Future Work

Our evaluation was performed on an English-language corpus and knowledge base. The Getty AAT contains translations of many records, facilitating an evaluation of entity linking pipelines in other languages. We leave this to future research.

A different and complementary approach to classify the objects in our dataset would be to make use of photos that are associated with them. While

photos are not available for the entire dataset and many of the available images do not capture the objects in detail, a possible improvement would use multimodal algorithms that combine images and text descriptions to link entities.

## Acknowledgements

The authors are grateful to the Data Intelligence Institute of Paris (diiP) for financial support and to TGIR Huma-Num (CNRS) for computational infrastructure support. Chaney thanks the European Research Council for financial support (ERC grant N°884847).

## References

- Omar Adjali, Romaric Besançon, Olivier Ferret, Herve Le Borgne, and Brigitte Grau. 2020. Multimodal entity linking for tweets. In *European Conference on Information Retrieval*, pages 463–478. Springer.
- Carmen Brando, Francesca Frontini, and Jean-Gabriel Ganascia. 2016. *REDEN: Named Entity Linking in Digital Literary Editions Using Linked Data Sets*. *Complex Systems Informatics and Modeling Quarterly*, (7):60 – 80.
- Mathias Coeckelbergs and Seth Van Hooland. 2021. Machine learning techniques for the management of digitised collections. In *Information and Knowledge Organisation in Digital Humanities*, pages 244–259. Routledge.
- Pierre-Philippe Combes, Laurent Gobillon, and Yanos Zylberberg. 2022. Urban economics in a historical perspective: Recovering data with machine learning. *Regional Science and Urban Economics*, 94:103711.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. *Autoregressive entity retrieval*. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Emmanuelle Delmas-Glass and Robert Sanderson. 2020. Fostering a community of pharos scholars through the adoption of open standards. *Art Libraries Journal*, 45(1):19–23.
- Leon Derczynski, Diana Maynard, Giuseppe Rizzo, Marieke Van Erp, Genevieve Gorrell, Raphaël Troncy,

- Johann Petrak, and Kalina Bontcheva. 2015. Analysis of named entity recognition and linking for tweets. *Information Processing & Management*, 51(2):32–49.
- Jennifer D’Souza, Anett Hoppe, Arthur Brack, Mohamad Yaser Jaradeh, Sören Auer, and Ralph Ewerth. 2020. The stem-ecr dataset: grounding scientific entity references in stem scholarly content to authoritative encyclopedic and lexicographic sources. *arXiv preprint arXiv:2003.01006*.
- James J Feigenbaum. 2016. A machine learning approach to census record linking. *Retrieved March*, 28:2016.
- Matthew Gentzkow, Bryan Kelly, and Matt Taddy. 2019. Text as data. *Journal of Economic Literature*, 57(3):535–74.
- Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldrige, Eugene Ie, and Diego Garcia-Olano. 2019. Learning dense representations for entity retrieval. *arXiv preprint arXiv:1909.10506*.
- Mark M Hall and David Walsh. 2021. Exploring digital cultural heritage through browsing. In *Information and Knowledge Organisation in Digital Humanities*, pages 244–259. Routledge.
- Stephen Hansen, Michael McMahon, and Andrea Prat. 2018. Transparency and deliberation within the fomic: a computational linguistics approach. *The Quarterly Journal of Economics*, 133(2):801–870.
- Matthew Honnibal, Ines Montani, Matthew Honnibal, Henning Peters, Sofie Van Landeghem, Maxim Samsonov, Jim Geovedi, Jim Regan, György Orosz, Søren Lind Kristiansen, Paul O’Leary McCann, Duygu Altinok, Roman, Grégory Howard, Sam Bozek, Explosion Bot, Mark Amery, Wannaphong Phatthiyaphaibun, Leif Uwe Vogelsang, Björn Böing, Pradeep Kumar Tippa, jeannefukumaru, GregDubbin, Vadim Mazaev, Ramanan Balakrishnan, Jens Dahl Møllerhøj, wvseeker, Magnus Burton, thomasO, and Avadh Patel. 2019. *explosion/spaCy: v2.1.7: Improved evaluation, better language factories and bug fixes*.
- J. Paul Getty Trust. 2022a. About the AAT. <https://www.getty.edu/research/tools/vocabularies/aat/about.html>.
- J. Paul Getty Trust. 2022b. Getty vocabularies OpenRefine reconciliation service. <https://www.getty.edu/research/tools/vocabularies/obtain/openrefine.html>.
- Jeffrey L Jensen, Daniel Karell, Cole Tanigawa-Lau, Nizar Habash, Mai Oudah, and Dhia Fairus Shofia Fani. 2022. Language models in sociological research: An application to classifying large administrative data and measuring religiosity. *Sociological Methodology*, 52(1):30–52.
- Craig A Knoblock, Pedro Szekely, Eleanor Fink, Duane Degler, David Newbury, Robert Sanderson, Kate Blanch, Sara Snyder, Nilay Chheda, Nimesh Jain, et al. 2017. Lessons learned in building linked data for the american art collaborative. In *International Semantic Web Conference*, pages 263–279. Springer.
- Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. End-to-end neural entity linking. *arXiv preprint arXiv:1808.07699*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Belinda Z. Li, Sewon Min, Srinivasan Iyer, Yashar Mehdad, and Wen-tau Yih. 2020. Efficient one-pass end-to-end entity linking for questions. In *EMNLP*.
- Elvys Linhares Pontes, Luis Adrián Cabrera-Diego, Jose G Moreno, Emanuela Boros, Ahmed Hamdi, Antoine Doucet, Nicolas Sidere, and Mickaël Coustaty. 2022. Melhissa: a multilingual entity linking architecture for historical press articles. *International Journal on Digital Libraries*, 23(2):133–160.
- Fangyu Liu, Ivan Vulić, Anna Korhonen, and Nigel Collier. 2021. Learning domain-specialised representations for cross-lingual biomedical entity linking. *arXiv preprint arXiv:2105.14398*.
- Pablo N Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. 2011. Dbpedia spotlight: shedding light on the web of documents. In *Proceedings of the 7th international conference on semantic systems*, pages 1–8.
- Sunil Mohan and Donghui Li. 2019. Medmentions: A large biomedical corpus annotated with umls concepts. *arXiv preprint arXiv:1902.09476*.
- Hannes Mueller and Christopher Rauh. 2018. Reading between the lines: Prediction of political violence using newspaper text. *American Political Science Review*, 112(2):358–375.
- Gary Munnely and Séamus Lawless. 2018. [Constructing a knowledge base for entity linking on irish cultural heritage collections](#). *Procedia Computer Science*, 137:199–210. Proceedings of the 14th International Conference on Semantic Systems 10th – 13th of September 2018 Vienna, Austria.
- GP Murdock, CS Ford, AE Hudson, R Kennedy, LW Simmons, and JWB Whiting. 2008. Outline of cultural materials (6th revise). *Human Relations Area Files*.

- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Mia Ridge. 2014. *Crowdsourcing our cultural heritage*. Ashgate Publishing, Ltd.
- Zejiang Shen, Kaixuan Zhang, and Melissa Dell. 2020. A large dataset of historical japanese documents with complex layouts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 548–549.
- Daniel Simig, Fabio Petroni, Pouya Yanki, Kashyap Papat, Christina Du, Sebastian Riedel, and Majid Yazdani. 2022. Open vocabulary extreme classification using generative models. *arXiv preprint arXiv:2205.05812*.
- Smithsonian Institution. 2022. Match Getty AAT. <https://github.com/Smithsonian/Match-Getty-AAT>.
- Tornike Tsereteli, Yavuz Selim Kartal, Simone Paolo Ponzetto, Andrea Zielinski, Kai Eckert, and Philipp Mayr. 2022. Overview of the sv-ident 2022 shared task on survey variable identification in social science publications. *arXiv preprint arXiv:2209.09062*.
- Seth Van Hooland, Max De Wilde, Ruben Verborgh, Thomas Steiner, and Rik Van de Walle. 2015. Exploring entity recognition and disambiguation for cultural heritage collections. *Digital Scholarship in the Humanities*, 30(2):262–279.
- Seth Van Hooland, Ruben Verborgh, Max De Wilde, Johannes Hercher, Erik Mannens, and Rik Van de Walle. 2013. Evaluating the success of vocabulary reconciliation for cultural heritage collections. *Journal of the American Society for Information Science and Technology*, 64(3):464–479.
- Shikhar Vashishth, Rishabh Joshi, Denis Newman-Griffis, Ritam Dutt, and Carolyn Rose. 2020. [MedType: Improving Medical Entity Linking with Semantic Type Prediction](#). *arXiv e-prints*, page arXiv:2005.00460.
- Albert Weichselbraun, Roger Waldvogel, Andreas Fraefel, Alexander van Schie, and Philipp Kuntschik. 2022. Slot filling for extracting reskilling and upskilling options from the web. In *International Conference on Applications of Natural Language to Information Systems*, pages 279–290. Springer.
- John Wilkerson and Andreu Casas. 2017. Large-scale computerized text analysis in political science: Opportunities and challenges. *Annual Review of Political Science*, 20(1):529–544.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. Zero-shot entity linking with dense entity retrieval. In *EMNLP*.

## A Appendix

### A.1 Data Sampling

There are significant imbalances in the frequency of words and expressions in the raw data that various museums provide. A significant part of this data consists of single word descriptions, like "paintings", "coins" and "figure". In order to have more information-rich data to annotate, we sampled a different number of examples of different types of fields independently. Our ARTIFACT dataset consists of:

- 50% Descriptions,
- 30% Titles,
- 20% Materials and techniques.

An effect of this sampling method is that for instance the presence of the "description" field for object A in our dataset does not mean that "title" filed of the same object will also be annotated with EL labels.

Another important characteristic of the data is that, since it comes from museum collections, objects are more numerous for some time periods than others. To have a more diversified sample, respecting the proportions mentioned above, we uniformly sampled descriptions from time windows, that is, we randomly selected the same number of strings from every two hundred years window after the year zero and from every five hundred years window before the year zero till 3000 b.c.

### A.2 Dataset statistics

Single-annotated data were generated once inter-annotator agreement rates were stable between previous annotation sessions. Additionally, for the single-annotation datasets, an additional data cleaning check was performed to filter or correct manual annotation errors related to label span detection. Table 7 shows the general statistics for the composition of the dataset.

Figure 3 shows the distribution of the length of the text (number of words) for object description and the

Data type	Annotations	US	UE	UM
Single-annotated	7,510	1,728	1,500	1,802
Cross-annotated	552	127	296	315

Table 7: Data composition by annotation type: Unique strings (US), Unique Entities (UE), and Unique mentions (UM).

distribution of the number of entities per text for the whole dataset. On average, the descriptions in our dataset are between 14 and 15 words long and have four entities.

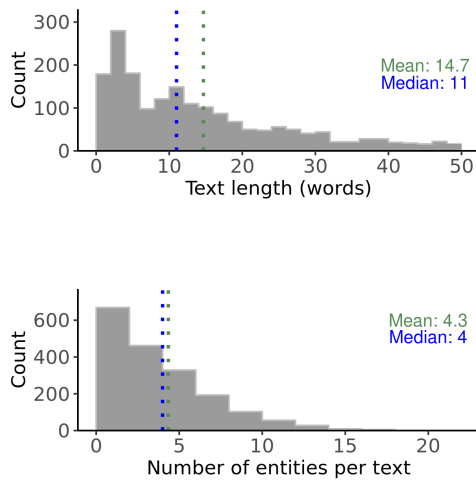


Figure 3: Top panel: Distribution of text length (words). Bottom Panel: Distribution of entities per text.

Table 3 shows the size of the text descriptions, the number of entities per text, and the number of unique entities for training, validation, and test sets, respectively.

The most common elements in our facets of interest are objects and materials. Table 8 shows the entity counts per facet, and unique identity counts per facet for each data split.

The top-10 most frequent entities by facet are provided in figures 4, 5, and 6.

Facet	Metric	Dataset		
		Train	Validation	Test
Activities	C	1,107	132	431
	UE	207	66	66
Materials	C	1,658	180	699
	UE	309	84	99
Objects	C	3,830	433	1,316
	UE	795	219	201
Styles & Periods	C	159	11	43
	UE	98	9	13

Table 8: Entity counts per facet (C), and unique entity counts per facet (UE) for training, validation, and test sets.

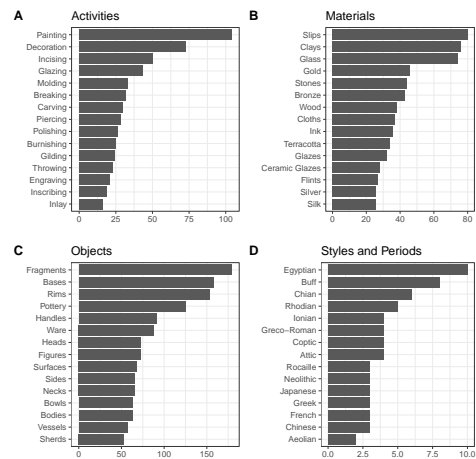


Figure 4: Top-10 entities by facets and counts for the training set

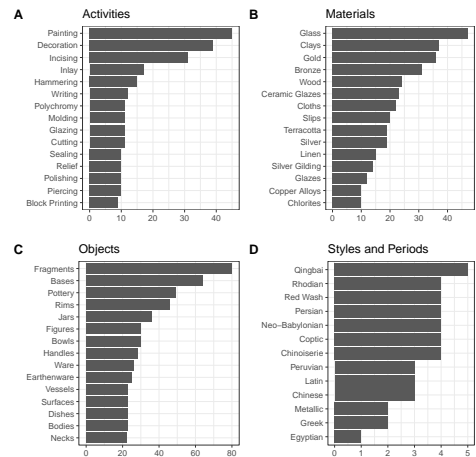


Figure 5: Top-10 entities by facets and counts for the test set

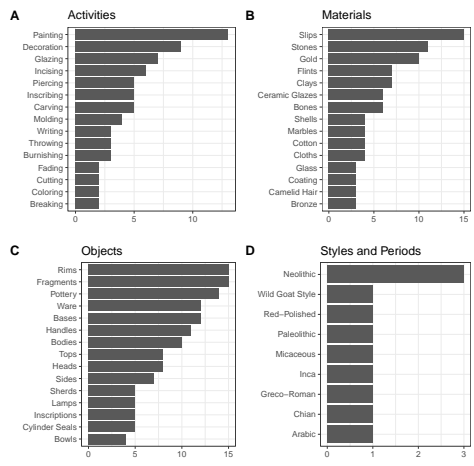


Figure 6: Top-10 entities by facets and counts for the validation set