

Enhancing Entity Alignment Between Wikidata and ArtGraph using LLMs

Anna Sofia Lippolis¹, Antonis Klironomos², Daniela F. Milon-Flores³, Heng Zheng⁴, Alexane Jouglar⁵, Ebrahim Norouzi⁶ and Aidan Hogan⁷

¹University of Bologna and ISTC-CNR, Italy

²Bosch Center for Artificial Intelligence (BCAI) | University of Mannheim, Germany

³Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, Grenoble, France

⁴School of Information Sciences, University of Illinois Urbana-Champaign, United States

⁵Computer science faculty, Université de Namur, Belgium

⁶FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, Germany

⁷Department of Computer Science, University of Chile, Chile

Abstract

Knowledge graphs (KGs) are used in a wide variety of applications, including within the cultural heritage domain. An important prerequisite of such applications is the quality and completeness of the data. Using a single KG might not be enough to fulfill this requirement. The absence of connections between KGs complicates taking advantage of the complementary data they can provide. This paper focuses on the Wikidata and *ArtGraph* KGs, which exhibit gaps in content that can be filled by enriching one with data from the other. Entity alignment can help to combine data from KGs by connecting entities that refer to the same real-world entities. However, entity alignment in art-domain knowledge graphs remains under-explored. In the pursuit of entity alignment between *ArtGraph* and Wikidata, a hybrid approach is proposed. The first part, which we call *WES* (Wikidata Entity Search), utilizes traditional Wikidata SPARQL queries and is followed by a supplementary sequence-to-sequence large language model (LLM) pipeline that we denote as *pArtLink*. The combined approach successfully aligned artworks and artists, with *WES* identifying entities for 14,982 artworks and 2,029 artists, and *pArtLink* further aligning 76 additional artists, thus enhancing the alignment process beyond *WES*' capabilities.

Keywords

Entity alignment, Wikidata, ArtGraph, Knowledge-graphs, Large Language Models

1. Introduction

Knowledge graphs (KGs) help organize and analyze information. In particular, cultural heritage KGs play a vital role in depicting and comprehending the significant relationships between individuals, objects, customs, events, and legacies, illuminating the tapestry of our collective heritage. Recently, such KGs have been used in several applications that aim to automatically analyze art [1, 2], investigate the connections between visual elements [3], represent traditional crafts [4], and improve the explainability of neuro-symbolic methods [5].

Rich information plays a vital role in applications related to cultural heritage. For example, discerning the emotional response elicited by an image in viewers is challenging due to the

SWODCH '23: International Workshop on Semantic Web and Ontology Design for Cultural Heritage, November 6 OR 7, 2023, Athens, Greece



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

intricate and subjective nature of human emotions and the abstract nature of visual arts [1]. To achieve this recognition of emotions, comprehensive background knowledge is necessary. A single KG that is either general-purpose or domain-specific might not be sufficient. Multiple KGs are often relevant to an application, but entity links identifying which nodes in these KGs refer to the same real-world entity are essential to integrate and leverage their data fully. Entity alignment is fundamental for integrating knowledge between different KGs by identifying those missing entity links. It ensures that the same real-world entities from the input graphs are aligned, even when using different identifiers in each graph. While entity alignment is widespread, its application in art-related KGs has yet to be explored.

In this work, we focus on computing entity alignments for two KGs: (1) Wikidata [6], which is a well-known community-driven KG representing general knowledge, and (2) *ArtGraph* [2], which is a recently proposed KG in the art domain. Multiple information gaps in these two KGs could be filled by identifying the KG nodes that refer to the same real-world entities. For instance, among the artworks mapped from *ArtGraph*, all of which possess an image, 5,356 on Wikidata lack the P18 property `image`, crucial for describing artworks. Another example is the P136 `genre` property, which is complete for most *ArtGraph* entities, but 8,830 artwork entities are without such a value on Wikidata. Other information that can enhance art-related data on Wikidata can be information about an artist or an artwork belonging to a movement or identifying art patrons as related to artists. Another notable observation is the disparity in information between WikiArt, from which *ArtGraph* is derived, and Wikidata. For instance, while WikiArt identifies Mario Cesariny as an artist, Wikidata highlights his role as a writer (ID: Q1360242). Linking both KGs would offer a more complete picture of the artists and artworks they describe.

Thus, there is an opportunity to fill information gaps and solve inconsistencies in both of the mentioned KGs to enable more robust applications. In this paper, we study entity alignment between Wikidata and *ArtGraph*. In particular, we aim to identify missing entity links between Wikidata and *ArtGraph*, focusing on the artists and artworks (Figure 1).

We propose a hybrid approach to perform entity alignment between *ArtGraph* and Wikidata. The first part of the approach, named *WES* (Wikidata Entity Search), uses traditional techniques such as Wikidata SPARQL queries (similar to [7]). This method is followed by *pArtLink*: a supplementary pipeline that uses sequence-to-sequence large language models (LLMs), such as Llama 2 [8], to enhance the results of the previous method. For aligning the artworks, we used *WES*, and for the artists, we first executed *WES* and then *pArtLink* to process artists that were not aligned using *WES*. We manually evaluate the correctness of the results. Using *WES*, we retrieved Wikidata entities for 14,982 artworks and 2,029 artists in *ArtGraph*. With *pArtLink*, we found Wikidata entities for 76 additional artists that *WES* left unaligned. We provide the code, the data, and the results in a public GitHub repository ¹.

The outline of the paper is the following: Section 2 introduces the background; Section 3 presents our proposed approach; Section 4 includes the results and evaluation of the approach; Section 5 discusses the approach; Section 6 concludes the paper².

¹<https://github.com/AntonisKl/Entity-Alignment-for-Art.git>

²Author contributions: A.K. and A.S.L. conceived, implemented, and tested the *pArtLink* and *WES* methods, respectively. D.F.M-F. worked on the reliability of the evaluation agreement and the respective implementation. E.N. and A.S.L. conducted data exploration. A.J., H.Z. conducted the literature review and contributed to the paper

2. Background

Knowledge graphs have been widely used to describe artworks and other cultural heritage artifacts, as discussed by Baroncini et al. [9]. For instance, ArCo, an Italian cultural heritage knowledge graph, collects 169 million triples describing 820 thousand cultural entities [10]. ArCo provides links to around 18.7K entities in external Linked Open Data (LOD) datasets, such as DBpedia, Wikidata, and Geonames. Europeana [11] is another key project for collecting and structuring data relating to cultural heritage across Europe. Works have looked at ways to link Europeana with external knowledge bases, such as WikiArt (see Task 4.3.1 in [12]). Other art-related KGs include Zeri for photos that document artworks [13], Nomisma for numismatic data (<https://nomisma.org/>), and PaintKG describing paintings for Chinese audiences [14].

ArtGraph is a KG that encapsulates concepts related to works of art. It was populated using WikiArt and DBpedia starting with the most popular genres and styles in WikiArt. *ArtGraph* is made up of 135,038 resources, including artworks, artists, genres, styles, movements, and more. In contrast with the aforementioned art-focused graphs, the information of *ArtGraph* is neither bound to a specific geographic region nor restricted by a specific type of visual artwork. In addition to its broad scope, WikiArt, which *ArtGraph* is based on, is built and maintained by Ukrainian developers [15], where taking into account the ongoing situation in the country, we are interested in supporting the project by connecting this KG to an external one, namely Wikidata.

ArtGraph includes Wikipedia URLs for the artists, but not for artworks. However, at the time of writing this paper, we noticed that there were artists who were missing a Wikipedia page URL (e.g. Ivan Vladimirov) and others that had an invalid one (e.g. Dana Levin). So we opted to use the Wikipedia URLs only as a first step of two in *WES*, as a basis to find a first batch of artists, and then execute *pArtLink* as an attempt to align the artists with missing or invalid Wikipedia URLs.

Regarding the entity alignment task, a survey by Zeng et al. [16] summarizes the relevant methods developed in recent years. Entity alignment has been applied in diverse scenarios, such as for geographic knowledge bases [17], for large-scale multilingual knowledge bases (DBpedia) [18], and for electric power marketing [19]. To the best of our knowledge, the alignment of entities within KGs related to the domain of art has not been investigated.

Our goal is to provide a hybrid solution for aligning entities between *ArtGraph* and Wikidata. We present the steps of our proposed solution in the following section.

3. Proposed approach

Our proposed approach uses two methods: *WES* and *pArtLink*. *WES* was executed and evaluated first for aligning artists and artworks. Afterwards, *pArtLink* was used to improve the results of *WES* in terms of artists' alignment. Thus, these two methods are used in a complementary manner. Before testing our approach, we examined the KGs (i.e., Wikidata and *ArtGraph*) to estimate their quality and extracted the data we needed as input for our methods. Both the data and the internals of our two methods are presented in this section.

writing. A.H. supervised the research. All authors provided critical feedback and helped shape the research.

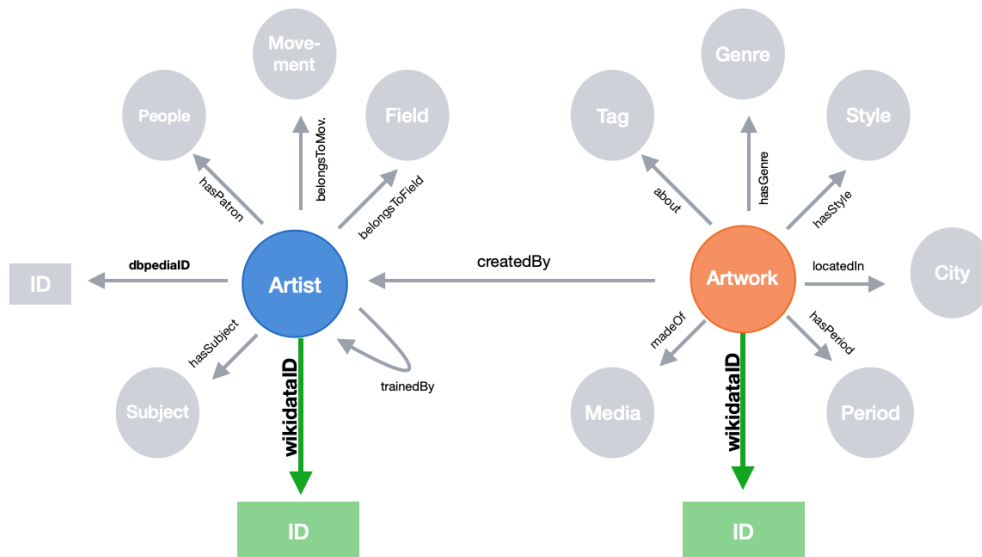


Figure 1: Artgraph ontology. The green arrows show the attributes added through this work.

3.1. Analysis of data resources

Building on the foundation laid out previously, this section explores the data resources integral to our approach. Central to this is the *ArtGraph* KG, which describes 135,038 resources and is accessible on Zenodo. Following our initial review of the *ArtGraph* KG, we conducted a content coverage analysis and identified any missing information.

Initial observations. Our preliminary observations revealed several issues. For instance, there is a discrepancy between the published KG and the ontology, specifically regarding the `belongtofield` attribute being used in one place and `belongstofield` in another. It is also necessary to consider the presence of alternative names or titles for the artworks concerning Wikidata. For example, alternative names may include middle names or alternative (translated) titles such as “Ezechiel’s Vision” versus “Vision of Ezechiel” or “Poppy Field” versus “Field with Poppies”. SPARQL queries were created to explore these issues to understand the extent of specific gaps concerning properties related to the entities of interest.

Artists exploration. Several information gaps were discovered for artists. The gender property is missing for 57 artists. While this may seem small, understanding gender can provide insights into the representation and diversity of artists in the dataset. Also, birth and death details are crucial for historical context. Yet, `birth_date` is missing for 20 artists, `death_date` for 425, `birth_place` for 1,341, and `death_place` for 1,566. Only 144 artists have information on `hasPatron`. This could be indicative of the challenges in tracing patronage historically. Similarly, high numbers of missing values in `trainedBy` (2,460), `belongsToField` (1,917), `relatedToSchool` (2,170), and `belongsToMovement` (1,850) suggest gaps in the professional and educational backgrounds of the artists.

Artworks exploration. Regarding artworks, similarly, numerous gaps in information were identified. The location of artworks is a significant gap, with `locatedIn` missing for 96,943 pieces. This could hinder efforts to trace the provenance or current location of the artworks. Furthermore, the vast majority, 108,153 artworks, are missing the `partOf` attribute, which

could provide insights into collections or series they belong to. It’s worth noting that some attributes like name, createdBy, hasGenre, image_url, and hasStyle for artworks are fully populated, suggesting that while there are significant gaps in some areas, others are well-documented.

To conduct entity aligning experiments, we extracted the names of the artists and their artworks in a 2-column CSV file, which was given as input for the proposed methods described in the following section.

3.2. Methodology

This section will introduce a two-step methodology designed to align artworks and artists on *ArtGraph* with the corresponding entities on Wikidata. For the artworks, we used *WES* (the first step of our approach) as a stand-alone method. For artists, we integrate our two strategies (i.e., *WES* and *pArtLink*) aiming to enhance both the accuracy and coverage of links. Given the limitations of Wikidata Entity Search and the absence of contextual knowledge in the input data, incorporating a complementary method based on state-of-the-art research would allow for identifying additional links and enable comparison with *WES*.

3.2.1. Method 1 – *WES*: Wikidata Entity Search

The first method, called *WES* for short, uses the names of artists and the titles of their artworks from *ArtGraph* to search for the corresponding entities in Wikidata. By matching entity labels between the two KGs, *WES* also provided a hint on the quality of Wikidata in terms of population completeness according to the definition proposed by [20]. This method was used to align both artworks and artists, as explained in this section.

Entity retrieval and linking to Wikidata for artworks. First, a SPARQL query using the *SPARQLWrapper* library in Python is executed to align the artwork’s title and creator with the corresponding title, description, and/or creator data present on Wikidata. This query employs EntitySearch, a practical method underpinned by ElasticSearch that is used to perform full-text searches on Wikidata pages, even in image captions. We made the word “The” optional before the title of an artwork to enhance the query. This ensures it was impossible to miss any titles that might not include it. As a limitation, it is not possible to make too many requests simultaneously, which is challenging for this case, as there is a large amount of data to process. To circumvent this, we have introduced a timeout mechanism. The methodology yields results for over 14,000 categorized artworks from an initial query set of more than 100,000 artworks.

Entity retrieval and linking to Wikidata for artists. Within the *ArtGraph*, hyperlinks point to artist profiles on Wikipedia. Given that Wikipedia entries are linked to from Wikidata, we have extracted these Wikipedia URLs to pinpoint the precise artist page title. This step is useful in avoiding disambiguation, as the Wikipedia page will point to the right entity in case of homonymy. Subsequently, the Python library *pywikibot* was employed to derive the corresponding Wikidata item. There were cases of artists that could not be identified, perhaps because of missing Wikipedia pages or invalid Wikipedia URLs (as discussed in Section 2). So, *Spacy OpenTapioca*³ was used on the remaining artists’ names directly.

³<https://spacy.io/universe/project/spacyopentapioca>




Figure 2: *pArtLink* architecture diagram.

After using this method for linking the artists, we observed that 423 of the 2,452 total artists were not linked. So, we proceeded towards a second complementary method using Natural Language Processing (NLP) and LLMs to fill the gaps for artists that were not linked.

3.2.2. Method 2 — *pArtLink*: Pipeline for entity retrieval and linking using LLMs

The second method, which we call *pArtLink*, complements *WES* with NLP techniques for mapping the artist names to Wikidata entity IDs after mapping them to the most likely Wikipedia article titles. The motivation behind developing this method was to find links for the more complex cases of artists missed by *WES*, for whom little context is often provided, i.e., the artists’ names and titles of their artworks. Our idea was to leverage modern LLMs’ retrieval and generative capabilities to enrich the context available for entity alignment.

The main processes involved are prompt generation, textual sequence generation, and entity identification and linking. The pipeline takes as input a CSV file with the artists and their artworks and outputs a CSV file with the artists and their Wikidata IDs. The workflow can be divided into the following modules that are arranged sequentially (Figure 2).

Prompt creation. This initial module of the pipeline creates one prompt for each artist and one of his/her artwork (arbitrarily chosen), which will then be given as input to an LLM. The prompt aims to make the LLM produce informative sentences for the artist. Another goal of the prompt is to yield answers with a specific format in order to be parsed later in an automated way. Furthermore, the answer from the LLM should not contain any additional redundant text (e.g., introductory remarks that rephrase parts of the prompt). With these goals as requirements, the constructed prompt template is: *You are a helpful AI assistant for finding information about artists. Your answer must contain only one line. Your answer must follow this template: “artist_name: info_about_artist”. Your answer must not include any other text than the answer itself. Write a short sentence about the artist {artist} who created {artwork}.*

Sentence generation. This module is responsible for generating for each artist a natural language sentence that describes them. It accomplishes this by feeding the prompts created by the previous module to a sequence-to-sequence LLM called *Llama 2* [8]. For the purposes of this work, a pre-trained version of the model called *Llama-2-Chat* with 7 billion parameters was used via a Python interface provided by *GPT4ALL* project [21]. *Llama-2-Chat* models are open-source models developed by Meta specifically for dialogue use cases. In comparison to open-source chat models, according to the original paper, these fine-tuned LLMs perform better on most benchmarks and are on par with some popular closed-source models like ChatGPT and PaLM in terms of helpfulness and safety, as per the human evaluations they conducted.

Entity retrieval and linking to Wikipedia. This pipeline step aims to identify the entities representing the artists in the given sentences and link them to Wikipedia. Specifically, for each sentence, this module retrieves the title of the Wikipedia article that corresponds to the artist mentioned in that sentence. If there is more than one possible title, all of them are retrieved for disambiguation purposes. This process is achieved by utilizing a sequence-to-sequence language model called *GENRE* (Generative ENTITY REtrieval) [22]. In this work, we used an instance of the pre-trained model on End-to-End Entity Linking from sentences to Wikipedia. The *GENRE* system employs a fine-tuned BART architecture to perform entity retrieval, connecting input text to distinct entities. It generates a unique entity name from the input text using constrained beam search to maintain the validity of the generated identifiers.

Wikipedia-to-Wikidata mapping. The final module of the workflow has the role of finding the correspondence between Wikipedia article titles and Wikidata entity IDs. A Python library called *Wikimapper*⁴ was used for performing this mapping. This tool creates an index of mappings from a Wikipedia SQL dump to accomplish this. Before the execution of the pipeline, a recent dump of the English Wikipedia was retrieved, and then the index was generated from it by *Wikimapper*. The coding interface of *GENRE* provides a way of specifying a custom Python method that functions as a converter from the Wikipedia article title to an ID. In this way, *Wikimapper*'s mapping functionality was injected in *GENRE* to perform the mapping internally before *GENRE* produces its output.

4. Results and evaluation

The results were separated depending on the category of the entities, i.e., artists and artworks (Figure 1). A sample of entities from each category was evaluated manually by the authors of this paper. Later, by using the Fleiss Kappa measure, we calculated the agreement of the evaluations for the result samples that were collectively judged.

4.1. Artworks

Artworks were linked using the *WES* method, generating 14,982 links to Wikidata. We then evaluated 100 artwork links randomly chosen from the results and concluded that 99 of these links were correct.

During the evaluation, an artwork was considered correctly linked if there was a convergence between the source WikiArt entity and the target Wikidata entity with regard to the title, author, image, and date properties. If all the mentioned fields except for the title matched, a slightly different title was accepted.

During the evaluation, we observed that if two or more artworks by the same author have a similar title, such as *The Drinkers* and *Drinkers in a Tavern* by Adriaen van Ostade, especially in the case one title is a subset of the other, then if one of them isn't listed on Wikidata, the other will often be incorrectly displayed as a result. These are particularly challenging cases.

We did not use our *pArtLink* method to link any artwork because we found some caveats that might influence with a greater chance the performance of *pArtLink* for artworks, than artists. In

⁴<https://github.com/jcklie/wikimapper>

particular, the state of Wikipedia in which *GENRE* was pre-trained (several years ago) is not the same as the current one. Given the higher number of artworks than artists, it is more probable that new artwork pages have been created on Wikipedia since then. Also, entity linking models are often trained for people, places, and organizations, and thus may not be suitable for titles of works, which may be verbose and contain common words like “Drinker” that are difficult to disambiguate. For these reasons, *GENRE* might not be able to perform entity linking for artworks with the same accuracy as for artists, though it would be interesting to explore this in future work.

4.2. Artists

Regarding the artists, our *WES* method was executed first to extract links for 2,452 artists. Wikidata entity IDs were found for 2,029 artists, out of which 100 were manually evaluated by the authors of this work; all were found to be correctly linked to Wikidata.

However, using *WES* led to 423 artists being left unlinked. We saw this as an opportunity to use our *pArtLink* method with the prospect that the abilities of LLMs will prove fruitful. In order to test *pArtLink*, we used as input 700 artists in total: the 423 unlinked by *WES* and 277 of the rest. This method found Wikidata entity IDs for 332 artists. We manually checked the correctness of these entities. Out of these artists, 117 were left unlinked by *WES* (i.e., 215 were also linked by *WES*) and 252 were correctly linked to Wikidata. Also, 76 of the correctly linked artists were artists that *WES* could not link.

Each entity link that was evaluated was considered to be correct if the following conditions were met: The retrieved Wikidata entity represents a person that has a name, nationality, and profession that match with those in the artist’s WikiArt page (searching by the artist’s artwork).

During evaluation, we detected some cases of erroneous and incomplete entity linking. Some artists, like Erol Akyavaş, do not exist in the English Wikipedia. Since *Wikimapper* is based on the English version of Wikipedia as a source for the mapping to Wikidata, such artists were either not linked to Wikidata or incorrectly linked (25 out of the evaluated 332). Other artists may have names that are common among different Wikidata entities (e.g. James Weeks). *GENRE* is used to disambiguate using the produced titles. Except for 19 manually disambiguated artists, a single Wikidata entity ID was produced by *pArtLink* for each of the others. Also, for 7 of the 332 evaluated artists, the LLM-generated sentence described another person with a similar name. As a result, *pArtLink* produced the wrong Wikidata entity ID. One example is Joaquim Rodrigo (painter), for which the LLM answered with a sentence describing Joaquín Rodrigo (music composer).

4.3. Reliability of evaluation agreement

After compiling the results of our two methods, *WES* and *pArtLink*, we performed a manual evaluation to measure the level of agreement between reviewers. To do this, we created a sample of 100 random artwork/artist entities. The task was to confirm whether the links to their Wikidata ID were correct. The evaluation was limited to two variables: ‘correct (1)’ or ‘incorrect (0)’. Five reviewers individually examined the same sample to test agreement among reviewers. An example of the sample file generated can be seen in Table 2.

Table 1

Entity linking results. Section 4.1 mentions why *pArtLink* was not used for artworks.

Entities	Methods	Total	Linked	Evaluated	Correctly linked	Incorrectly linked
Artworks	<i>WES</i>	116475	14982	100	99	1
	<i>pArtLink</i>	-	-	-	-	-
Artists	<i>WES</i>	2452	2029	100	100	0
	<i>pArtLink</i>	700	332	332	252	76

Table 2

An example of the generated sample to evaluate the alignment of the entities (artist and artwork). Six reviewers analyzed 100 entities and their respective Wikidata entity IDs.

Artist	Artwork	Wikidata entity ID	Reviewer 1	Reviewer 2	...	Reviewer 6
Eleanor Fortescue-Brickdale	An Outcast	Q22034080	1	1	...	1
Ivan Aivazovsky	Calm Sea	Q105071676	0	0	...	1
...
Lesser Ury	Café Baue	Q62116516	0	0	...	0

Our process began by examining the artwork samples. After the first round of evaluations, we decided to apply the well-known Fleiss Kappa (FK) measure to assess the level of agreement among the reviewers. Specifically, FK allows evaluating the ‘inter-rater reliability’ of nominal data (e.g., correct/incorrect) in situations with 3 or more raters [23]. Predefined values allow interpreting the score obtained with FK. For example, if the score is less than 0, the agreement is poor; if it is between 0.01 and 0.2, the agreement is slight; between 0.61 and 0.8, it is a substantial agreement; and if the score is between 0.8 and 1, the agreement is almost perfect. The resulting FK for the first evaluation of artwork samples was 0.011, indicating a ‘slight agreement’. After some discussion among the reviewers, it was discovered that most of the disagreements were due to special cases on the Wikidata page: cases where the language of the artwork title in the sample (e.g., *La Pergola*) diverged from that of the Wikidata page (e.g., *After lunch*), cases where the artist’s name was abbreviated or slightly modified, e.g., Ivan Aivazovsky (sample) and Iwan Aiwasowski (Wikidata), and situations where the entity label was exclusively within the caption of the illustration (e.g., Q62116516). Therefore, we decided to perform a second round of evaluations for the artwork sample after defining the criteria mentioned in sections 4.1 and 4.2. For the second round, the obtained FK was 0.018, indicating a ‘slight agreement’. However, despite the low concordance, we estimated that more than 90% of the entities in the sample were marked as correct by the majority of the reviewers. During our research, we found that a limitation of the FK score is that it is affected by the no-variation/bias in the reviewed data. In situations like this, it is important to incorporate other measures to get a more comprehensive view of the agreement. In the study of [24], the authors suggested using the Percentage of Agreement (PA) measure alongside FK. Similarly, in the work of [25], the researchers introduced a modified version of the FK measure called Free Marginal Multirater Kappa (K_{free}) to address issues related to prevalence and unbalanced classes. By employing both approaches, we computed the PA and K_{free} values for our artwork sample, resulting in

the following outcomes: a PA of 95% and a K_{free} of 0.9. This time both results were higher than the obtained FK measure.

Subsequently, we extended our evaluation to the artist entities. For this case, we performed only one round of evaluation with the same pre-defined criteria as in Subsection 4.2, i.e., if the name of the artist, its occupation, and nationality in the Wikidata ID corresponded to the one on the sample, regardless of the languages and slight modifications, the sample was marked as correct. As anticipated, there was a distinction between the artist and the artwork evaluation due to the utilization of different methods. Consequently, the FK value of 0.88 indicated a ‘substantial agreement’ reflecting less prevalence in the data. The other metrics also exhibited consistency with the FK score, with PA remaining at 95% and K_{free} at 0.89.

To conclude with the evaluation of concordance, we can confirm that low FK values do not always reflect inter-rater reliability when the sample is unbalanced towards a particular label. Conversely, we can observe that both scores of artwork and artist, calculated with the K_{free} and PA metrics, are high. This result allows us to verify that for the reviewers, the entities were correctly associated with their Wikidata ID in most of the cases.

5. Limitations

In this work, we discuss a framework to fill information gaps and inconsistencies about artists and artworks for two knowledge bases, comprised of two methods: the *WES* method and the *pArtLink* method. These approaches, while successfully aligning artworks and artists, encounter some limitations. In the case of artworks with similar names for the same artists, the *WES* method may find false positives. Moreover, the results of the *pArtLink* method heavily rely on the quality of the sentences that get generated by an LLM. It could be that some artists were not present in the training dataset of the LLM. In this case, the LLM may not be able to produce an informative enough sentence so that *GENRE* can successfully link the entities mentioned in the text to Wikipedia. Furthermore, to the best of our knowledge, there is no perfect metric to measure agreement among reviewers. Readers should be aware that using the FK metric alone may result in low scores due to unbalanced assessments. A feasible solution is to incorporate complementary metrics for a more complete analysis and broaden the assessment categories beyond correct or incorrect.

6. Conclusion and future work

In this paper, we introduced two methods for entity alignment and applied them to align entities between the *ArtGraph* and Wikidata KGs. On the one hand, *WES* uses traditional querying techniques. On the other hand, *pArtLink* uses recently developed LLMs. Our evaluations demonstrate that these methods are complementary and can yield accurate results. With the entities aligned in this work, the information gaps of both *ArtGraph* and Wikidata can be filled by completing missing data in a subsequent endeavor.

In this research, we investigated the task of entity alignment in art-related knowledge bases. However, there are still avenues for further exploration and enhancement. One of the immediate steps forward in terms of practical impact is to establish a collaboration with Wikidata by

uploading the evaluated and corrected datasets. Additionally, to validate the versatility and scalability of our framework, it would be beneficial to test its efficacy on other art-related knowledge bases. Finally, we plan to expand our *pArtLink* method on artworks, and to link entities with Wikipedia articles in languages other than English.

Acknowledgments Heng Zheng thanks the grants, United States Institute of Museum and Library Services RE-250162-OLS-21, Alfred P. Sloan Foundation G-2022-19409, the University of Illinois Urbana Champaign, and the University of Groningen for supporting me in participating in the 2023 International Semantic Web Summer School which led to this contribution. Antonis Klironomos is partially affiliated with the EU project Graph Massiviser (GA 101093202). Anna Sofia Lippolis is supported through the WHOW project (EU CEF programme - grant agreement no. INEA/CEF/ICT/A2019/2063229). Ebrahim Norouzi thanks the Information Service Engineering group at FIZ Karlsruhe for supporting me in participating in the International Semantic Web Summer School 2023. Daniela F. Milon-Flores is part of the TRACES project (<http://traces-anr-fns.imag.fr/>) funded by ANR: Agence Nationale de la Recherche (France) and FNS: Fonds National Suisse (Swiss). Alexane Jouglar is part of the ARIAC project (<https://www.digitalwallonia.be/fr/publications/trail/>).

References

- [1] S. Aslan, G. Castellano, V. Digeno, G. Migailo, R. Scaringi, G. Vessio, Recognizing the emotions evoked by artworks through visual features and knowledge graph-embeddings, in: P. L. Mazzeo, E. Frontoni, S. Sclaroff, C. Distanto (Eds.), *Image Analysis and Processing. ICIAP 2022 Workshops*, Springer International Publishing, 2022, pp. 129–140.
- [2] G. Castellano, V. Digeno, G. Sansaro, G. Vessio, Leveraging Knowledge Graphs and Deep Learning for automatic art analysis, *Knowledge-Based Systems* 248 (2022) 108859.
- [3] A. Kouretsis, I. Varlamis, L. Limniati, M. Pergantis, A. Giannakouloupoulos, Mapping art to a knowledge graph: Using data for exploring the relations among visual objects in renaissance art, *Future Internet* 14 (7) (2022) 206.
- [4] N. Partarakis, V. Doulgeraki, E. Karuzaki, G. Galanakis, X. Zabulis, C. Meghini, V. Bartalesi, D. Metilli, A web-based platform for traditional craft documentation, *Multimodal Technologies and Interaction* 6 (5) (2022) 37.
- [5] N. Díaz-Rodríguez, A. Lamas, J. Sanchez, G. Franchi, I. Donadello, S. Tabik, D. Filliat, P. Cruz, R. Montes, F. Herrera, Explainable neural-symbolic learning (x-nesyl) methodology to fuse deep learning representations with expert knowledge graphs: The monumai cultural heritage use case, *Information Fusion* 79 (2022) 58–83.
- [6] D. Vrandečić, M. Krötzsch, Wikidata: a free collaborative knowledgebase, *Commun. ACM* 57 (10) (2014) 78–85.
- [7] A. Waagmeester, E. L. Willighagen, A. I. Su, M. Kutmon, J. E. L. Gayo, D. Fernández-Álvarez, Q. Groom, P. J. Schaap, L. M. Verhagen, J. J. Koehorst, A protocol for adding knowledge to wikidata: aligning resources on human coronaviruses, *BMC Biology* 19 (1) (2021) 12.
- [8] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al., Llama 2: Open foundation and fine-tuned chat models, *arXiv preprint arXiv:2307.09288* (2023).
- [9] S. Baroncini, B. Sartini, M. Van Erp, F. Tomasi, A. Gangemi, Is dc:subject enough? a

landscape on iconography and iconology statements of knowledge graphs in the semantic web, *Journal of Documentation* 79 (7) (2023) 115–136.

- [10] V. A. Carriero, A. Gangemi, M. L. Mancinelli, L. Marinucci, A. G. Nuzzolese, V. Presutti, C. Veninata, Arco: The italian cultural heritage knowledge graph, in: C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois, F. Gandon (Eds.), *The Semantic Web – ISWC 2019*, Springer International Publishing, Cham, 2019, pp. 36–52.
- [11] A. Isaac, B. Haslhofer, Europeana linked open data –data.europeana.eu, *Semantic Web* 4 (3) (2013) 291–297.
- [12] M. Lefferts, C. Concordia, L. Anastasiou, A. Jahnke, M. Kittelmann, H. Manguinhas, Europeana cloud - deliverable 4.3 - a report and a plan on future directions for improving metadata in the europeana cloud, Tech. rep., Project report, Europeana Cloud, Deliverable D4.3, 2016 (2016).
- [13] M. Daquino, F. Mambelli, S. Peroni, F. Tomasi, F. Vitali, Enhancing semantic expressivity in the cultural heritage domain: Exposing the zeri photo archive as linked open data, *J. Comput. Cult. Herit.* 10 (4) (jul 2017).
- [14] H. Wu, S. Y. Liu, W. Zheng, Y. Yang, H. Gao, Paintkg: the painting knowledge graph using bilstm-crf, in: *2020 International Conference on Information Science and Education (ICISE-IE)*, 2020, pp. 412–417.
- [15] Wikiart visual encyclopaedia blocked in russia, <https://shorturl.at/cxCU8> (4 2022).
- [16] K. Zeng, C. Li, L. Hou, J. Li, L. Feng, A comprehensive survey of entity alignment for knowledge graphs, *AI Open* 2 (2021) 1–13.
- [17] K. Sun, Y. Zhu, J. Song, Progress and challenges on entity alignment of geographic knowledge bases, *ISPRS International Journal of Geo-Information* 8 (2) (2019).
- [18] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, C. Bizer, DBpedia –a large-scale, multilingual knowledge base extracted from wikipedia, *Semantic Web* 6 (2) (2015) 167–195.
- [19] W. Meng, D. Zhang, T. Guo, Z. Zong, Y. Liu, Y. Wang, J. Li, W. Zhu, Research on the typical application of knowledge graph in power marketing, in: *2021 International Conference on Computer Engineering and Application (ICCEA)*, 2021, pp. 318–321.
- [20] C. Brando, N. Abadie, F. Frontini, Linked data quality for domain-specific named-entity linking, in: *Conférence Extraction et Gestion des Connaissances (EGC) Atelier Qualité des Données du Web (QLOD)*, Reims, France, 2016, hal-0239918.
- [21] Y. Anand, Z. Nussbaum, B. Duderstadt, B. Schmidt, A. Mulyar, Gpt4all: Training an assistant-style chatbot with large scale data distillation from gpt-3.5-turbo, <https://github.com/nomic-ai/gpt4all> (2023).
- [22] N. De Cao, G. Izacard, S. Riedel, F. Petroni, Autoregressive entity retrieval, in: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, OpenReview.net, 2021.
- [23] J. L. Fleiss, Measuring nominal scale agreement among many raters., *Psychological bulletin* 76 (5) (1971) 378.
- [24] M. L. McHugh, Interrater reliability: the kappa statistic, *Biochemia medica* 22 (3) (2012) 276–282.
- [25] J. J. Randolph, Free-marginal multirater kappa (multirater k [free]): An alternative to fleiss' fixed-marginal multirater kappa., Online submission (2005).