



Deployment and Maintenance of Europeana DSI core services - SMART 2016/1019

CONTRACT NUMBER - 30 - CE - 0885387/00-80

DELIVERABLE

C.2 Usage pattern report 2

Revision	1.0
Date of submission	30 April 2018 (M8)
Author(s)	Dasha Moskalenko, Europeana Foundation, Ash Marriott, Europeana Foundation, Antoine Isaac, Europeana Foundation, Monica Paramita, University of Sheffield, Paul Clough, University of Sheffield
Dissemination Level	Public



Co-financed by the European Union
Connecting Europe Facility

REVISION HISTORY AND STATEMENT OF ORIGINALITY

Revision History

Revision No.	Date	Author	Organisation	Description
0.1	19-4-2018	Dasha Moskalenko, Ash Marriott, Antoine Isaac	Europeana Foundation	First version
0.2	26-4-2018	Dasha Moskalenko, Ash Marriott, Antoine Isaac, Monica Paramita, Paul Clough	Europeana Foundation University of Sheffield	Final version
1.0	30-4-2018	Victor-Jan Vos, Julia Schellenberg	Europeana Foundation	Final review

Statement of originality:

This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.

The sole responsibility of this publication lies with the author. The European Union is not responsible for any use that may be made of the information contained therein.

'Europeana DSI is co-financed by the European Union's Connecting Europe Facility'.

TABLE OF CONTENTS

Executive summary	4
Introduction	5
1. Development of UGC component for Europeana Collections	6
1.1 Existing solution	6
1.2 Scoping development	6
1.3 New UGC platform	7
1.4 Future iterations	7
2. Search improvements and evaluation	8
2.1 Measuring performance for Europeana search	8
2.2 Evaluation of the entity-based autocompletion	11
2.2.1 Heuristic evaluation of auto-complete functionalities	11
2.2.2 System-oriented evaluation of performance and accuracy	13
2.2.3 Coverage of the Entity Collection knowledge graph	16
2.3 Main lessons learnt and identified progress points	17
3. Europeana Pro usage patterns	19
3.1 Conclusion	20
Appendix A: evaluation of the coverage of the Entity Collection knowledge graph	21
A.1 Completeness of information	21
A.1.1 Number of populated properties per entity type	21
A.1.2 How do the properties in the Entity Collection compare to other Web sources?	22
A.2 Multilingual availability	24
A.3 Proportion of CH entities (from other Linked Data sources) that are available in the EC	28
Appendix B: evaluation of linkages between Europeana objects and the Entity Collection	29

Executive summary

Research and analysis of usage patterns led to three main areas of work to improve access to Europeana:

1. How can we capitalize on the success of the Europeana 1914-1918 campaign in other themes?

We discovered that one of the main success factors of the campaign was the engagement of the general public in sourcing material/artifacts (User generated content) from the First World War. Therefore, we aimed to make this activity part of other campaigns (e.g. Migration) as well. The existing solution for uploading UGC was created several years ago to address a generic use-case for cultural heritage institutions, and was not suited to the simpler needs of Europeana. Therefore, we developed a new solution more robust and flexible than the former solution.

2. How can we serve users content that is most relevant to their search query?

Major efforts go into the further development of the search functionality with an emphasis of exploring the effectiveness and completeness of the search-related algorithms. This aims making it clearer where improvements can be made to enhance user experience on Europeana. Besides updating our search server infrastructure and investigating extension of the search function we started exploring the evaluation of the (Entity-based) auto-complete component. The evaluation provided insights on how the auto-complete API supports different types of user searches. At the moment, the auto-complete API provides limited features in handling imperfect queries, e.g., those with misspelling or different word order. However, given that users provided the correct queries, the auto-complete API were able to suggest relevant entities with good success rate for all entity types. Further, we evaluated the accuracy of the auto-complete feature. The results show that the auto-complete API achieves a good success rate in retrieving the correct entities for users once users typed 5 characters or more. In general, the results show that the success rate of the auto-complete API is good, given that the entities users queried for exist in the EC. If the entities relating to user queries do not exist, however, the auto-complete API will not be able to provide correct suggestions.

3. Did we improve the user satisfaction of Europeana Pro?

We launched the new Europeana Pro website, main communication tool towards our professional target markets in September 2017. We set some propositions (e.g. ease of use, navigation) to evaluate the effects of the redesign. The results of a user survey concluded that we meet the propositions set for the redesign. Thus, we consider the redesign of the website to be successful.

Introduction

This document describes how Europeana determines where and how users experience difficulties with the use of the Platform and the actions we take to overcome such difficulties with systematic improvements over time. The report refers to the user-facing parts of the Platform, Europeana Pro (for all professional markets) and Europeana Collections (for end-users).

For both products, we defined objectives in December 2017 derived from a key performance indicator of Europeana DSI-3, or an already known user issue, as indicated by user research or Commission Services. In general, development of a new feature or a design change is based on either or both of these.

Our most common methodologies to collect user feedback are: surveys, heatmaps, nps scores, and feedback buttons. We also continuously experiment with other methodologies on how to best evaluate where users are experiencing difficulties with the user-facing products and thus be able to improve the service.

This deliverable should be read in combination with Deliverable C.3 Data Access Pattern Report (updated January 2017 and April 2018)¹, which describes how users access and engage with Europeana data.

¹ The reports can be found in the section project documentation on the Europeana DSI-3 project page. View at <https://pro.europeana.eu/project/europeana-dsi-3>

1. Development of UGC component for Europeana Collections

Over the past three months we developed a solution to collect user generated content (UGC) which is suited for use on any of the thematic collections. The reason that we prioritized this development is due to the following research question: "How can we capitalize on the success of the Europeana 1914-1918 campaign in other themes?"

When we started to analyze the Europeana 1914-1918 campaign we discovered that one of the main success factors of this campaign was the engagement of the general public in sourcing material/artifacts (UGC) from the First World War. User generated content (UGC) is often very unique and unexpected. It also tells the story from the perspective of an ordinary citizens which resonates well with the general user of our site.

Primary stakeholder groups:

- Partner institutions assisting us in submitting the data
- Contributors of stories and objects to Europeana
- The wider community with an interest in the stories we are telling

1.1 Existing solution

The existing solution for uploading UGC was created several years ago to address a generic use-case for cultural heritage institutions to upload UGC. It provided a fully customisable experience, and was not suited to the simpler needs of Europeana.

End users and cataloguers gave us feedback that the interface of the Europeana 1914-1918 platform was unnecessarily heavy, with the need to register and log-in being a barrier to entry. They additionally stated that the format of the questionnaire was clunky and not easy to fill-in. In an interview scenario, they stated it was difficult to keep opening the different accordions and move between pages to fill in data.

For Europeana, the platform provided no direct mapping to EDM which meant that we were losing data in the ingestion process. The RunCoCo project² was no longer maintained making it difficult for us to continue to use.

1.2 Scoping development

Our partners dedicate time to collecting UGC stories, and 97% of our items and stories on Europeana 1914-1918 have historically come from partner collection days making them the

² The existing solution which will be replaced by the new UGC tool.

primary focus of our improvements. We wanted to make it simpler for users to submit UGC online; as well as simplifying the cataloguing process. To simplify the cataloguing process it was important that the data would be quick and easy to ingest, therefore it needed to be modelled accordingly.

Our hypothesis was that we can:

1. Reduce the barrier for the interviewer at collection days to enter the story of the user, reducing the time spent making the interview.
2. Obtain a constant flow of submissions without the need to log-in and edit them afterwards.
3. Gain more Public Domain content by allowing additional licensing options on upload.
4. Get more views of content by the use of linked data at source.

1.3 New UGC platform

The new solution that we developed is more robust and flexible than the former solution we had. The key improvements are:

- To make it easier for our partners to interact with the system we shortened the instruction guide from 20 pages to 4 pages. We were able to do that by simplifying the process for data upload.
- To make it easier for contributors we reduced the barrier to submission through removal of a login/authentication system and made the new UGC platform mobile accessible.
- To make the UGC data more accessible for the general user, we use linked data in both locations and topics which improves connectivity between records and consistency in search terms. Therefore this should improve the overall findability of UGC content.

1.4 Future iterations

We believe that users would still benefit from being able to enrich their story as they discover more about themselves, and find more objects. This requires a login functionality that would allow users to re-access and edit their content. We will monitor user interactions and feedback on the new platform to identify how to further develop it to increase its value. In Europeana DSI-4 we will continue the development on this platform.

2. Search improvements and evaluation

In this report we also look at our activities to improve our search and discovery functionality. Specifically exploring the effectiveness and completeness of the search-related algorithms.

2.1 Measuring performance for Europeana search

Major efforts go into the further development of the search functionality. We have seen major improvements in Europeana DSI-1 and DSI-2, and under this project while Europeana continues to invest in improving the search functionality. Currently, there is an emphasis on how to measure the search functionality effectiveness and completeness, so that we can more objectively assess the results of our work.

Since Europeana DSI-2, Europeana has embarked on a long-haul effort of defining new, finer indicators for appraising search and discovery on Europeana (in Europeana DSI-2 there was only one KPI 'Percentage of all user searches matched to an entity from the Entity Database'). Previous search improvement reports developed during the Europeana DSI-1 and DSI-2 projects³ guide this ongoing Europeana DSI-3 work. The emphasis is on metrics immediately recoverable from our log data. That is to say, on data currently being harvested, rather than information that would require further software-engineering work to be done (other than processing of logs implemented earlier). The metrics are also connected to specific search and discovery components, making it clearer where improvements can be made to enhance user experience on Europeana.

The focus on search development in the last reporting period has been rather on updating our search server infrastructure and investigating extension of the search function so that it can accept both metadata and full-text documents, in order to meet the requirement of the coming Europeana Newspapers thematic collection. As a result we could not considerably extend the set of metrics we have reported on in the past period, while previous investigations gave us many more candidate indicators⁴. We could, however, start exploring the evaluation of the (Entity-based) auto-complete component and report on a basic metric for satisfaction with the recently deployed Entity Pages component.

The updated set of defined metrics of the search functionality is:

³ The most recent Search Improvement Report is DSI-2 D6.3 available at https://pro.europeana.eu/files/Europeana_Professional/Projects/Project_list/Europeana_DSI-2/Deliverables/d6.3-search-improvement-report.pdf

⁴ We have gathered a Europeana Information Retrieval Component Inventory (<https://docs.google.com/spreadsheets/d/1RAv1oZ3rVyfKC7bnwW29RYUW0sy-rnDZE6j0d2Yup5k>) and are defining new KPIs with the University of Sheffield, (https://docs.google.com/document/d/16TKUfpZVM7m3SXjgfpD1_9Z2QvScxrj8MIpdGHbCgb4/).

Component	Evaluation Criteria	Description	Value for the period of September to December 2017	Value for the period of January to April 2018
Search box	nDCG	Normalised Discounted Cumulative Gain: a measure of how high in the result list clicked items are found. See https://en.wikipedia.org/wiki/Discounted_cumulative_gain#Normalized_DCG . Used as a reporting metric in previous DSIs' Search Improvement Plan Progress Reports.	0.562	0.611
	Percentage of queries with clicked results	The proportion of queries that are followed by a click on their results, as opposed to queries without clicks (which are then assumed not to have brought relevant results). Note that this figure may be strongly depressed by internal operations checking datasets, when many searches are launched but few results viewed.	3.4%	16.3%
	Entity Collection coverage	Percentage of searches matchable by entities within the Entity Collection. Used as a reporting metric in previous DSIs' Search Improvement Plan Progress Reports (with a KPI of 30%)	34.2%, 45.7% of entity queries	58.7% (based on a smaller sample due to time constraints)
	nDCG upon completion of autosuggest term by user	nDCG as applied to the entity list supplied by the autocomplete. Available only for the testing phase of the Entity Collection API (by Europeana staff) so we expect it to be artificially higher than (coming) production results.	0.887 (unchanged until a script is implemented to register users' selection within the list of entities suggested by the auto-complete)	
Filters	Frequency of use	The number of times filters of search results are activated by users	50.1%	26.4%
	Frequency of item access after filter applied	The number of times the use of a filter is followed by a click on one of the items in the new (refined) search results.	31.4%	29.7%

Similar Items	Frequency of use	The number of times users have clicked on one of the items suggested as part of the 'similar items' listed for an object being browsed.	0.15%	0.24%
	nDCG	Normalised Discounted Cumulative Gain: a measure of how high the similar item clicked was in the list of all similar items suggested.	0.60	0.66
Auto-complete	See section 2.2			
Entity Pages	Bounce rate	For a detailed discussion on the bounce rate please see C.2 Usage pattern report (M4) ⁵	N/A (the component was not deployed for the full period)	61.76%

The activities to enhance performance of the search and discovery algorithms are continuing along the lines mentioned in the previous search improvement reports (i.e. changing weights of fields used for the ranking, considering completeness of metadata records and presence of thumbnails, etc.). EF also redesigns some interface components that will impact some metrics (such as the design of the search result page).

One example is the contribution of the autosuggest for search accuracy. In September 2017, we introduced an auto suggest query strings functionality for our users when entering a search term in the search box to help harmonize user queries and provide a more satisfying and reliable search experience. The functionality, now deployed for persons, places and concepts, appears on the main search bar on Europeana Collections, on Thematic Collections and on a smaller search bar on object pages. Our expectation is that having a query completed by the autosuggest will help users arrive faster at the items that are relevant for them, and we can measure this using the nDCG. The assumption is that clicking an item higher up the list of results signifies that this list was more useful.

There are two stages of development of the autosuggest feature:

- Stage 1: Autosuggest by keywords - This offers the completion of search queries using terms available from the Entities API. - This option is already developed.
- Stage 2: Autosuggest using Entities URI - This completes the search query, but searches using the Entity URI which covers items dereferenced to an entity, as well

⁵ The report can be accessed on the Europeana DSI-3 project page in the section project documentation. View at <https://pro.europeana.eu/project/europeana-dsi-3>

as the equivalent searches in different languages. - This option is scheduled to be developed.

Once the second stage of development is completed, we'll be able to do a comparison of two metrics:

- nDCG of results obtained when search uses a keyword;
- nDCG of results obtained when search uses an Entity URI.

We expect that results obtained using an Entity URI will be more relevant and users will find content quicker through these result sets.

As of April 2018, Stage 2 is not completed. However, we embarked on a finer-grained evaluation of the auto-complete function, seeking to evaluate whether the entities it returns are relevant to the user queries, even if these entities are not used directly in the search for objects. The results of this work are available in the following section.

2.2 Evaluation of the entity-based autocompletion

The auto-complete function based on Europeana's Entity Collection supports users in searching for items on Europeana. We evaluated three aspects that affect its performance.

Note: by April 2018, the rank of suggested entities in the Europeana Collections portal was different to the rank produced by the AutoComplete API. In this report, the following evaluation is based on the AutoComplete API, which is the common 'building block' to be used by the portal and other (third-party) services.

2.2.1 Heuristic evaluation of auto-complete functionalities

The first aspect is the functionality supported by the auto-complete component. Users' search queries often contain cases such as misspelling or different word order. We carried out a heuristic evaluation to identify whether these cases are supported by the auto-complete component.

In this evaluation, a small number of queries that represented different scenarios (e.g., misspelling, word order, etc.) were tested. We report how the auto-complete component supports these different scenarios in the following:

- *Misspellings*: not supported.
 - In the current version of the auto-complete API, misspelled or mistyped query terms are not supported by the API.
 - E.g., "leonarda" query does not retrieve "leonardo da vinci" entity.
- *Partial match*: not supported
 - All terms in user queries have to exist in the retrieved entity.

- E.g., if a user search for the query: “Johann Justin Preissler”, the auto-complete API will not retrieve the entity: “Johann Preissler” (although they are the same entity). Due to incomplete information of the entity’s full name in the EC, the ‘skos:prefLabel’ value is only a partial match to user’s query and therefore is not retrieved by the auto-complete API.
- *Different word order*: not supported
 - The auto-complete API requires the query words to contain the same order as the entity name.
 - E.g., “vinci leonardo” or “leonardo vinci” queries do not retrieve “leonardo da vinci” entity. These cases may be better supported once the alternative labels (‘skos:altLabel’) are taken into account by the auto-complete suggester (given that the different-word-order queries are available as ‘skos:altLabel’ values).
- *Query with symbols*: not supported
 - The logs included queries with symbols, such as double quotes (e.g., “Leonardo da Vinci”), tags (e.g., “who:(Peter+Paul+Rubens)”, “type:”Photo”).
 - These types of queries are not supported and will not be matched into the correct entities.
- *Are sufficient information given to users to help them choose the correct entity from the list of suggestions?*
 - *E.g., are the entity types provided?* No. At the moment, only the entity names (‘skos:prefLabel’ values) are shown to the users. The entity types are not shown to users.
 - *E.g., are there cases of ambiguous suggestions?* Yes, there are often cases of ambiguity due to the redundant information in the EC. E.g., searching for “London” finds two “City of London”, both inherited from the Geonames database.

These findings suggest that the auto-complete API currently provides limited features to support users in their searches. The performance of the API can be improved using the following recommendation:

- *Data integration from other sources*. Integration of more labels (‘skos_prefLabel’) derived from other sources can avoid the low recall issue due to incomplete information in the EC.
- *Removal of duplicate information*. At the moment, there are many duplicate entities found in the Places entities (inherited from Geonames). Removal of these duplicate information will reduce the ambiguity issues and improve the suggestion results.
- *Entity types should be made available to users*. This information will help users in choosing the correct entities in the suggestions.
- *Support for partial matching*. At the moment, partial matches are not supported in order to prioritise the results’ precision and reduce the noise. However, as a result, the auto-complete API produces a low recall for cases with “less-than-perfect” queries (e.g., queries with misspelling or different word order). A partial matching

should therefore be considered if performing an exact-match of queries do not return any results.

2.2.2 System-oriented evaluation of performance and accuracy

The second aspect we evaluated is the performance and accuracy of the auto-complete function. We have carried out system-oriented evaluation that reports the response time of the auto-complete API and the precision of its results.

On average, the response time is 286ms (tested using the [responsiveness-testing scripts](#) in GitHub across 54 entities). To evaluate the accuracy of the auto-complete features, we built a multilingual test set using Google Analytics query logs from January 2018-February 2018. Manual assessment was carried out to extract the 10 most popular entities for each EU language portal of the Europeana sites. For those languages where the data were insufficient, a longer period was used to extract the query logs (February 2017-February 2018).

A number of filters were applied into the dataset. Firstly, entities that did not represent the chosen language portal were filtered out. Secondly, entities that did not exist in the EC were also filtered out at this stage. For each language, the top two entities were selected for the multilingual dataset. For each entity, we extracted the entity ID by searching for entries in the EC that match the entity query. For cases where more than one entity was found, the entity with the highest likelihood is chosen. Only one ID can be associated for each entity in the dataset. We show the statistics of the resulting dataset⁶ in the following table:

Number of languages	24 languages
Number of queries	48 queries (2 queries per language)
Type of queries	16 Agents; 28 Places; 4 Concepts
Length of queries (in number of characters)	Mean=8.98 (SD=3.90), Min=4, Max=18

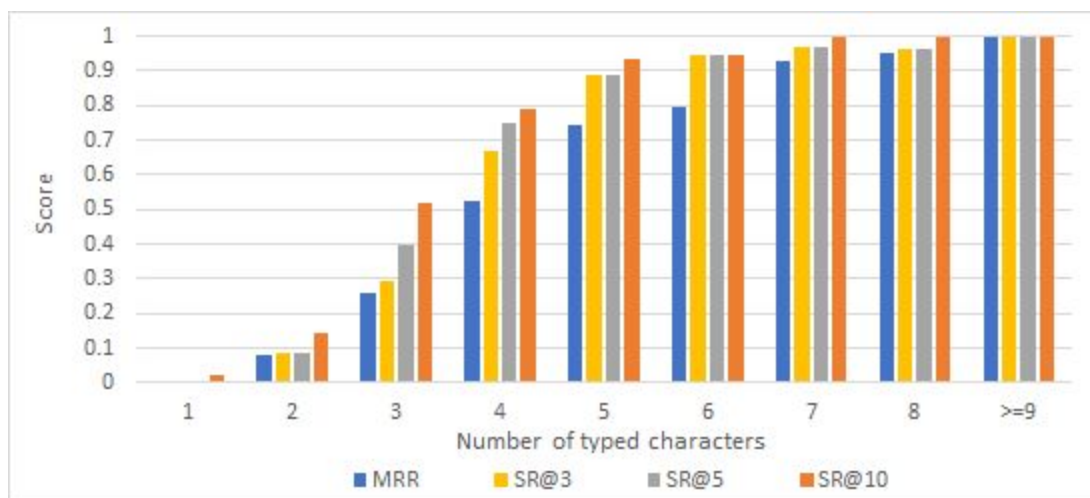
Statistics of the multilingual dataset

We recursively evaluated the performance of the auto-complete API using the first n character of each query (entity name). Since only one entity is deemed to be correct for each query, for each query, we evaluated the Reciprocal Rank (RR), which is the inverse of the rank of the correct entity in the results. E.g., if the correct entity is retrieved in rank 1: $RR=1$, rank 2: $RR=\frac{1}{2}$, rank 3: $RR=\frac{1}{3}$, etc. We reported the Mean Reciprocal Rank (MRR) over all the 48 queries in the dataset. We also evaluated the success rate at top- n (SR@ n), which

⁶

https://docs.google.com/spreadsheets/d/112n9mFDJ4JQFBVYFnqh3DZS_HaDxaSjHObHCwmUgGFk/edit#gid=667992447

measures the average percentage of queries that retrieve the correct entity in the top n results.⁷ The results are shown in the following figure.



Evaluation results of the auto-complete feature (all queries, n=48)

The results show that the auto-complete API, given the first 3 characters of user queries, was able to retrieve the correct entities in the top 10 suggestions in 50% cases (SR@10=0.52). Over 90% cases were completed (i.e., the correct entities were found by the auto-complete API in the top 10) after users have typed the first 5 characters (SR@10=0.93).

The performance of the auto-complete features for the different types of queries were also evaluated. The three following figures show the results for the “Agent” queries, “Place” queries and “Concept” queries, respectively.

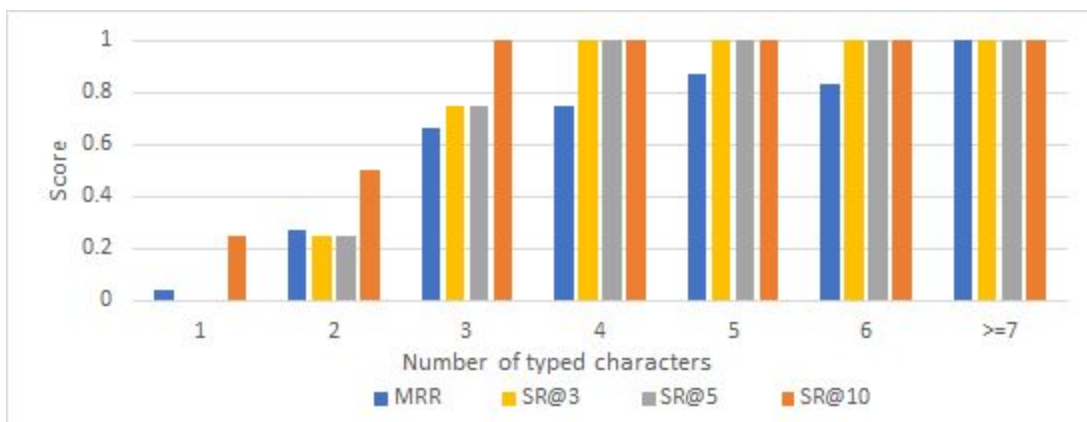


Evaluation results of the auto-complete feature (“Agent” queries, n=16)

⁷ More information about the metrics are shown in <https://staff.fnwi.uva.nl/m.derijke/wp-content/papercite-data/pdf/cai-survey-2016.pdf> (Section 5.2)



Evaluation results of the auto-complete feature ("Place" queries, n=28)



Evaluation results of the auto-complete feature ("Concept" queries, n=4)

The results show that the auto-complete API achieves a high SR@10 (above 0.8) for Agents, Places, and Concepts, after users have typed 5, 4 and 3 characters, respectively. In some cases, the correct entity was found but not suggested in the top rank, which resulted in the lower MRR score compared to SR@n scores. After further analysis, we found that these can be caused by two scenarios:

1. The same entity was included in the EC multiple times. E.g.:
 - "Izola" (a town in Slovenia) was found in two entities: <http://data.europeana.eu/place/base/85422> and <http://data.europeana.eu/place/base/85423>
 - "Sotkamo", a region in Finland, is found in two entities: <http://data.europeana.eu/place/base/86518> and <http://data.europeana.eu/place/base/86519>

Since the test dataset only link one user query to one EC entry, retrieving both entries in the suggestion results may cause the rank of the correct entity to appear to be lower.

2. Some user queries are ambiguous. E.g., user searching for "Palermo", may intend to search for the city (<http://data.europeana.eu/place/base/150586>), the province

(<http://data.europeana.eu/place/base/142989>), or Blinky Palermo, a German painter (<http://data.europeana.eu/agent/base/163139>). Since only one correct entity is predicted for each query in the gold standard, retrieving all these entities may result in the rank of the predicted entity appear to be lower. The predicted entity may also not accurately represent the user intent.

These results show that the auto-complete API achieves a good success rate in retrieving the correct entities for users once users typed 5 characters or more. There is room for improvement for the ranks of the correct entities in the suggestions results. One approach to investigate is by using the popularity of the entities (e.g., based on search frequency) in ranking the suggestions in order to improve the rank of frequently-searched entities⁸. Removal of duplicate entities will also improve the auto-complete results.

2.2.3 Coverage of the Entity Collection knowledge graph

In general, the automatic evaluation results show that the success rate of the auto-complete API is good, given that the entities users queried for exist in the EC. If the entities relating to user queries do not exist, however, the auto-complete API will not be able to provide correct suggestions. To understand better the latter issue, we have assessed the coverage of the data in the EC itself, its knowledge graph.

The auto-complete function indeed relies on the coverage and multilinguality of the EC knowledge graph. It is only able to suggest relevant entities to user queries if these entities exist there. Furthermore, the entities should also contain good multilingual availability (in the label attributes) in order to support the matching of multilingual queries.

As reported in Section 2.1, the "Entity Collection coverage" is 34.2%, up to 45.7% of our Entity queries.⁹ This is good for a start, but can probably be improved. Our assessment effort seeks to (1) gain better insight on the measures reported in earlier sections and (2) guide future efforts of curating the EC¹⁰. The details of this assessment can be found in Appendix A.

This assessment gives hints on how the auto-completion component can be enhanced. Improving the knowledge graph component of the EC is not the only relevant option to consider, though. Support of the EC to the search process also derives from the linkage

⁸ The current ranking algorithm uses some level of popularity, though it is based on Wikipedia (<http://people.aifb.kit.edu/ath/>) instead of Europeana.

⁹ Figures reported in Europeana DSI-2 D6.3 Search Improvement Report at https://pro.europeana.eu/files/Europeana_Professional/Projects/Project_list/Europeana_DSI-2/Deliverables/d6.3-search-improvement-report.pdf. The query testbed is available at <https://docs.google.com/spreadsheets/d/15xSicsev5v1JIK3an8JOXGvFWkM75G56lkBtVl-fNFkQ/>.

¹⁰ The curation process of the Entity Collection is described in: <https://docs.google.com/document/d/1A5Rb3Oe9edin5gdRpgFILLR0YPUodVOel3SdcBP00dA/>. The next update of the Entity Collection knowledge graph is foreseen for after the deployment of the new data ingestion system Metis.

done between objects and the EC entities in through automatic semantic enrichment or links contributed by providers¹¹. We have performed an initial assessment of this linkage, the details details of which can be found in Appendix B.

2.3 Main lessons learnt and identified progress points

In the previous section, we have evaluated the performance of the auto-complete API using a heuristic evaluation and a system-oriented evaluation. The heuristic evaluation provided insights on how the auto-complete API supports different types of user searches. At the moment, the auto-complete API provides limited features in handling imperfect queries, e.g., those with misspelling or different word order. However, given that users provided the correct queries, the auto-complete API were able to suggest relevant entities with good success rate for all entity types. Recommendations have been made to further suggest approaches to improve the performance of the API.

We also further reported that the performance of the auto-complete relies on the coverage and multilinguality of the Entity Collection itself. Therefore, we evaluated the intrinsic quality of the EC (reported in Appendix A). The availability of links between Europeana's Cultural Heritage Objects (CHOs) and relevant entities, which is an important aspect for supporting Europeana search in the future, was also evaluated (reported in Appendix B).

During this work, we have created three datasets to assist in evaluating the EC and the components it supports:

1. A dataset for evaluating the coverage of the EC¹²
2. A small exploratory dataset for evaluating the quality of the entity links¹³
3. A dataset for evaluating the auto-complete API¹⁴

We also identified future avenues to improve the EC for supporting Europeana search:

Increase EC coverage

The coverage evaluation results (in Appendix A) show that the EC has small coverages compared to entities available in Wikidata, ULAN and AAT. More analysis is needed to analyse the scope of the missing data, and, if necessary, to combine the data with other data sources. Another approach is to evaluate the coverage of the external data sources against Europeana queries. External data should be added if it manages to improve the coverage of Europeana queries (currently 45.7%).

¹¹ <https://pro.europeana.eu/page/europeana-semantic-enrichment>

¹² <https://basecamp.com/1768384/projects/5774755/messages/76013887>

¹³ <https://docs.google.com/spreadsheets/d/16VNWeIO0sHqi8pFYB4h00V6kpWW0rTMaeCd6ad1bGWg/>

¹⁴ https://docs.google.com/spreadsheets/d/112n9mFDJ4JQFBVYFnqh3DZS_HaDxaSjHObHCwmUgGFk/edit#gid=667992447

Other linked-data sources (e.g., entity card section in Google search¹⁵, RKD Museums¹⁶) should be considered as a gold standard to compare against. If the EC coverage against these sources is low, we may investigate the possibility of integrating their knowledge graphs into the EC.

Improving the quantity and quality of the CHO-Entities Linkages

The results (in Appendix B) currently show that the quantities of the CHO-Entities linkages can be further improved, especially for Places and Concepts; both are relevant for describing many Europeana CHOs. The low number of entity links to Agents also need to be further investigated, both in the context of increasing the quantity of linkages (to available entities) and also in alignment to the work in curating a suitable Entity Collection for Europeana CHOs (e.g., to improve the coverage of Agents EC). Future work should also include gathering a larger-scale annotations to create gold-standard to compare against. One specific point will be how to reprise the work of the earlier Europeana Network task force on evaluation and enrichments¹⁷. In particular, prior to launching a larger-scale annotation task, further work is also needed to create guidelines of ideal entities that Europeana would like to have. E.g.:

1. What are the ideal “Concept” entities for a CHO?
 - a. Is the type of the object enough (e.g., “painting”, “oil painting”)? Should it also contain general terms (e.g., “Arts”, or “Visual Arts” for paintings)? Should it also describe the contents of the CHOs (e.g., identifying the content of the painting, e.g., “lady”, “milkmaid”)?
 - b. Should the concepts include the name of the collections? E.g., should all items in the “World War I” collections have “World War I” as a concept?
 - c. Should black-and-white photos be linked to the “black-and-white photographs” entity¹⁸? If that was the case, should all colour photographs be linked to the “colour photographs” entity¹⁹?
2. What are the ideal “Place” entities for a CHO?
 - a. If a city is identified, e.g., “London”, should the parent entities be linked to (e.g., “England”, “United Kingdom”, “Great Britain”)? Please note that the relations between these entities are often already available in the EC KG itself.
 - b. If duplicate entities are found in the EC, should them all be linked to the CHO? E.g., there are three duplicate entities of “England, UK” in the EC.
 - c. Should the “Place” entities contain the provider’s place? Should it contain any related places, e.g., where it was created? Or location related to the subject of the painting?
3. What are the ideal “Agent” entities for a CHO?

¹⁵ E.g., <https://www.google.co.uk/search?q=leonardo+da+vinci>

¹⁶ E.g., <https://rkd.nl/en/explore/artists/Leonardo%20da%20Vinci>

¹⁷ <https://pro.europeana.eu/project/evaluation-and-enrichments>

¹⁸ <http://data.europeana.eu/concept/base/1676>

¹⁹ <http://data.europeana.eu/concept/base/1677>

- a. At this moment, only the creator of the CHO is identified as a link (given that the relevant entity exists in the EC).
- b. Should the subject of the paintings be included as the “Agent” (if exists in the EC)? Currently, there are no limitations of the entity types that can be included in any property. E.g., a painting subject (‘dc:subject’) can contain Place, Concept and Agent entities.

Improve auto-completion feature capabilities

This report identifies some limitations of the auto-completion feature, e.g., the lack of support for handling misspelling, different word order, etc. Providing these capabilities will help support users better in their tasks.

Carry out task-based evaluation of the entity-based knowledge cards and pages and auto-completion

Future work should aim to evaluate the usefulness of the entity-based knowledge cards and pages and the auto-completion in supporting user tasks. A task-based evaluation can be designed to investigate if the use of the knowledge cards can help users to complete their tasks more efficiently. Auto-completion features can be evaluated by comparing the search results (with the use of auto complete vs not using the auto-complete). I.e., are relevant results found in higher rank when the auto-complete is used? More metrics to capture user satisfaction should also be investigated in the future.

3. Europeana Pro usage patterns

Europeana Pro is Europeana’s main communication tool towards our professional target markets (Cultural Heritage Institutions, Academic Research, Education, and Creative Industries) as well as towards the wider Network of data partners and experts. We launched the new Europeana Pro website in September 2017.

Our propositions for the redesign of Pro were as follows:

- 27.8% of users rated the old Europeana Pro was easy to use with 9 and 10. We aim to increase this to 35%.
- 27.8% of users rated the ease of navigation of the old Europeana Pro with 9 and 10. We aim to increase this to 35%.
- 33.8% of users rated that old Europeana Pro met their expectations with 9 and 10. We aim to increase this to 35%.

A few months after the new site went live we asked our users the same three user satisfaction questions as we did on the old website. Here are the results.

- For the question “How easy is it to use the website?” 36,21% of our users scored the new Europeana Pro website at 9-10.

- For the question “How easy is the website to navigate?” 35,6% of our users scored the new Europeana Pro website at 9-10.
- For the question “Does the website meet your expectations?” 35.63% of our users scored the new Europeana Pro website at 9-10.

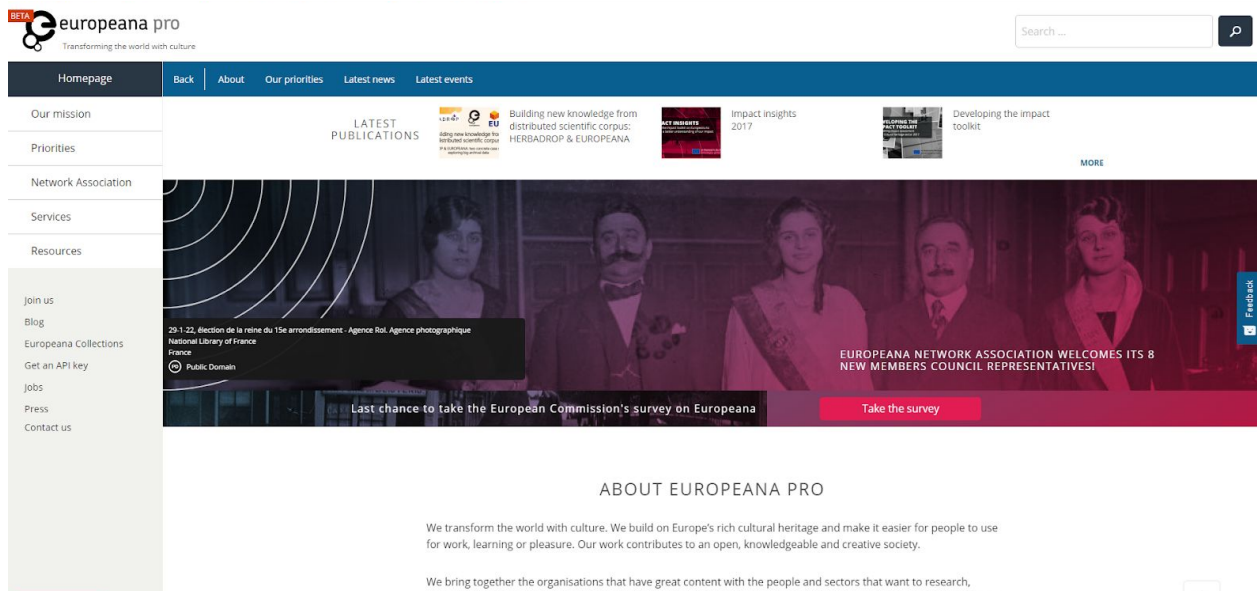
These results meet the propositions we set for the redesign of Europeana Pro. Thus, we consider the redesign of the website to be successful.

3.1 Conclusion

This year we will continue to improve the user satisfaction of Europeana Pro. In the coming months our primary focus will be the main navigation. Observing users interacting with Europeana Pro highlighted that they often neglect to use the main navigation and prefer to navigate through the website using the in page modules. We have a few hypothesis reasoning why the main navigation is not as popular as we expected:

- Users don't understand how the navigation functions.
- Users don't see the navigation.
- Users don't understand the labeling of the navigation.

Before making any design or functional changes to the website we will first conduct further research to validate any of the above hypotheses. The outcomes of this investigation will be reported in the next C.2 report (August 2018).



New Europeana Pro

Appendix A: evaluation of the coverage of the Entity Collection knowledge graph


A.1 Completeness of information

A.1.1 Number of populated properties per entity type


Entity Type	Total	Number of Populated Properties per Entity
Agents	165,005	Mean=29.68 (SD=12.18), Median=27, Min=15, Max=129
Places	215,802	Mean=20.08 (SD=6.00), Median=19, Min=17, Max=236
Concepts	1,572	Mean=39.65 (SD=26.84), Median=30, Min=16, Max=112

EC information completeness/richness


For each entity type, we selected a number of important properties and reported the number of entities that contain the selected properties.



Number of "Agent" records with specific properties (165,005 entities)



Number of "Place" records with specific properties (215,802 entities)



Number of "Concept" records with specific properties (1,572 entities)

A.1.2 How do the properties in the Entity Collection compare to other Web sources?

We compared the attributes for 8 different painters (manually selected) in the EC and Wikipedia infoboxes. In most cases, the EC (Table 2) contains 6 properties: "alternative name labels", "biographical summary", "the date of birth", "place of birth", "date of death", and "place of death". Meanwhile, Wikipedia infoboxes contain up to 13 attributes for some painters (shown in Table 3). The additional Wikipedia attributes include information about

their nationalities, notable works, movement/period, etc. Although these attributes are not included in the EC, this information is often included in the short summary described in the “biographical information” properties in the EC. Therefore, the EC coverage of information available for users is comparable to Wikipedia. Furthermore, this information can be extracted to create additional attributes to support a more advanced search (e.g., searching by ‘nationalities’ or ‘movement/period’) if necessary.

Comparison between the EC and the use of Wikipedia infoboxes may introduce some bias in the evaluation, due to the use of Wikipedia infoboxes in populating DBpedia (which is one of the resources for building the EC).

	Leonardo da Vinci	Vincent van Gogh	Edvard Munch	Johannes Vermeer	Caravaggio	Claude Monet	Frida Kahlo	Salvador Dali
Alt label (may contain birth name, alternative name, alternative spelling)*	✓	✓	✓	✓	✓	✓	✓	✓
Biographical information (summary)	✓	✓	✓	✓	✓	✓	✓	✓
Date of birth	✓	✓	✓	✓	✓	✓	✓	✓
Place of birth	✓	✓	✓	✓	✓	✓	✓	✓
Date of death	✓	✓	✓	✓	✓	✓	✓	✓
Place of death	✓	✓	✓	✓	✓	✓	✓	✓
Is related to**				✓				
<p>+ This property captures inter-agent relations, e.g., between student and teacher. Similar information is shown in the “Education” attribute in Wikipedia infobox (see Table 3).</p> <p>* These attributes are not shown in the entity page shown to the users.</p>								

Properties for 8 different painters in Europeana’s EC

	Leonardo da Vinci	Vincent van Gogh	Edvard Munch	Johannes Vermeer	Caravaggio	Claude Monet	Frida Kahlo	Salvador Dali
Birth name	✓	✓			✓	✓	✓	✓
Date of birth	✓	✓	✓	✓	✓	✓	✓	✓
Place of birth	✓	✓	✓	✓	✓	✓	✓	✓
Date of death	✓	✓	✓	✓	✓	✓	✓	✓
Place of death	✓	✓	✓	✓	✓	✓	✓	✓

Resting place		✓						✓
Nationality		✓	✓	✓		✓	✓	✓
Education		✓		✓			✓	✓
Known for (e.g., paintings, sculptures)	✓	✓	✓	✓	✓	✓	✓	✓
Works / Notable work	✓	✓	✓	✓	✓	✓	✓	✓
Movement (e.g., Baroque, Renaissance)	✓	✓	✓	✓	✓	✓	✓	✓
Patron(s)					✓	✓		
Spouse							✓	✓

Properties for 8 different painters in Wikipedia (infobox attributes)

A.2 Multilingual availability

In this experiment, we evaluate the multilingual availability of the ‘skos:prefLabel’ attribute for each entity type in the 24 EU official languages²⁰ (shown in Table 4).

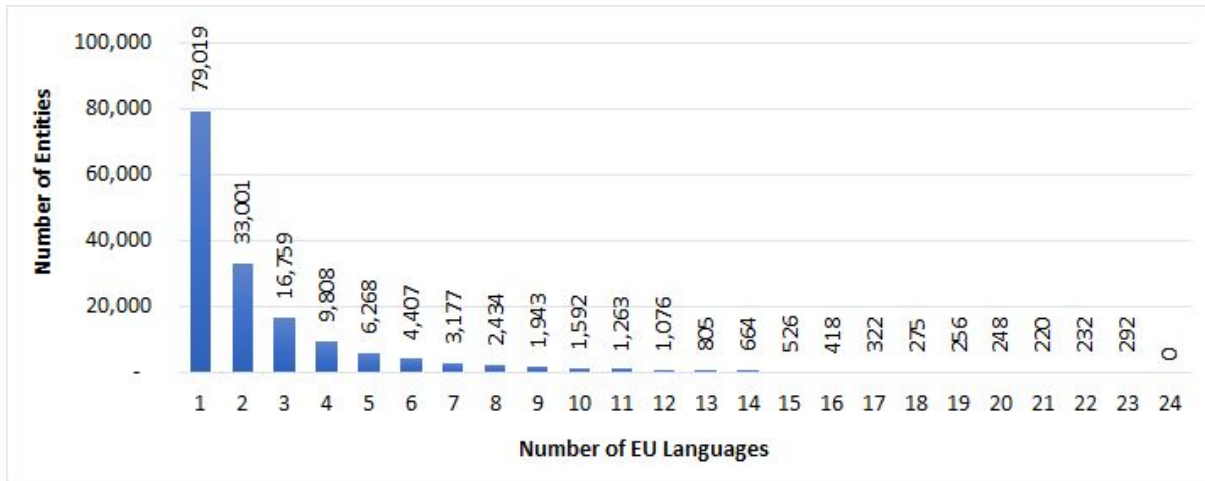
Bulgarian (BG)	Spanish (ES)	Hungarian (HU)	Polish (PL)
Czech (CS)	Estonian (ET)	Italian (IT)	Portuguese (PT)
Danish (DA)	Finnish (FI)	Lithuanian (LT)	Romanian (RO)
German (DE)	French (FR)	Latvian (LV)	Slovak (SK)
Greek (EL)	Irish (GA)	Maltese (MT)	Slovenian (SL)
English (EN)	Croatian (HR)	Dutch (NL)	Swedish (SV)

List of 24 official EU languages

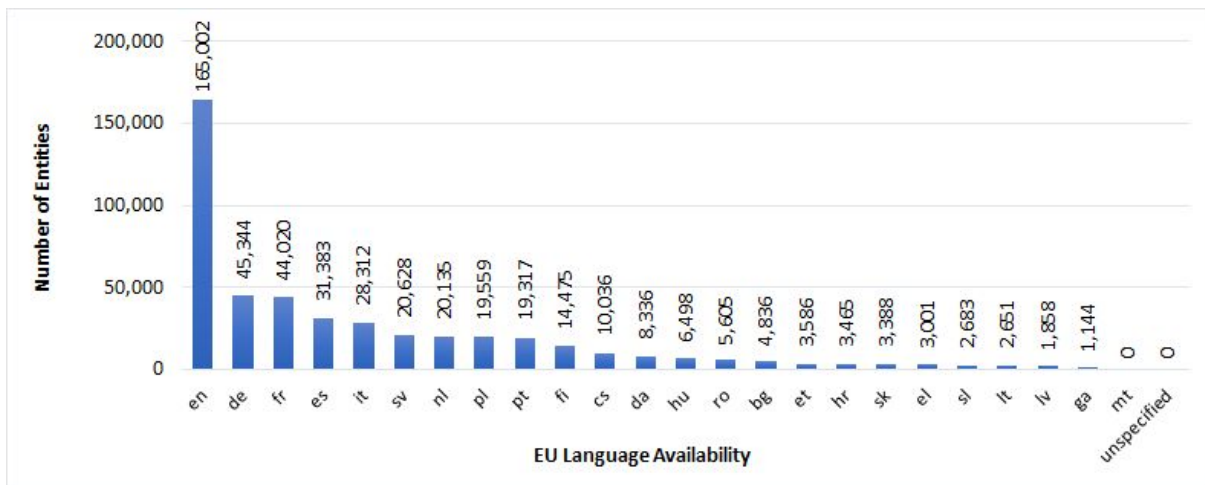
For each entity type, we evaluate the number of languages that were included in the property “skos:prefLabel” for each entity. This number represents the number of different languages that an entity name (or label) is available in. We also report the number of entities with values in the “skos:prefLabel” property for each EU language. The results for each entity type are shown below.

Agents: 165,005 entities

²⁰ https://europa.eu/european-union/topics/multilingualism_en



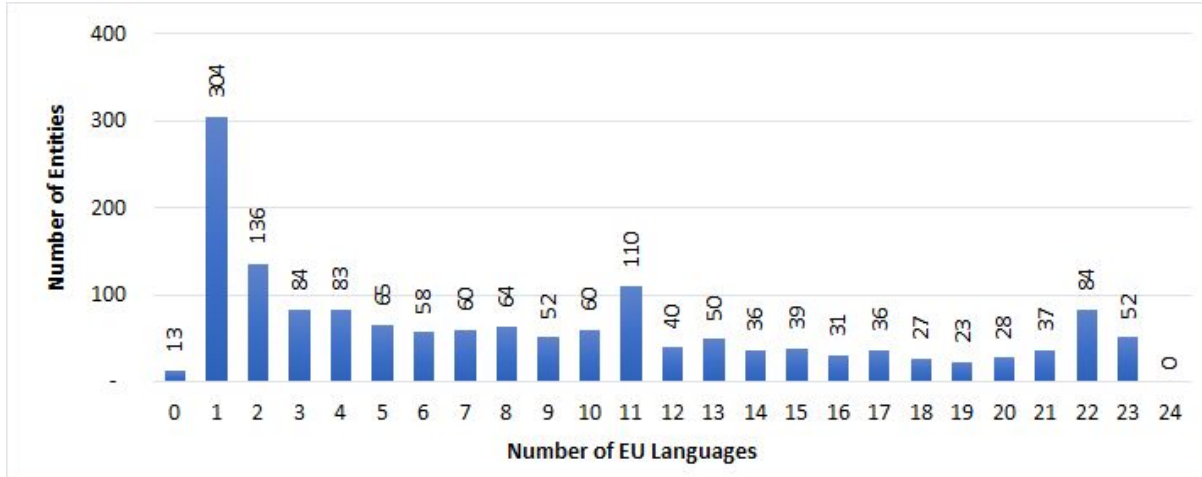
Multilingual availability (*skos:prefLabel*) for Agents (165,005 entities)



Multilingual availability (*skos:prefLabel*) for Agents (165,005 entities) per language

On average, each “Agent” entity contain 2.82 EU language (preferred) labels (SD=3.2, min=1, max=23, median=2). Almost all but 3 entities of the “Agents” type have English labels. The remaining three entities have either Italian, German, Spanish and French labels.

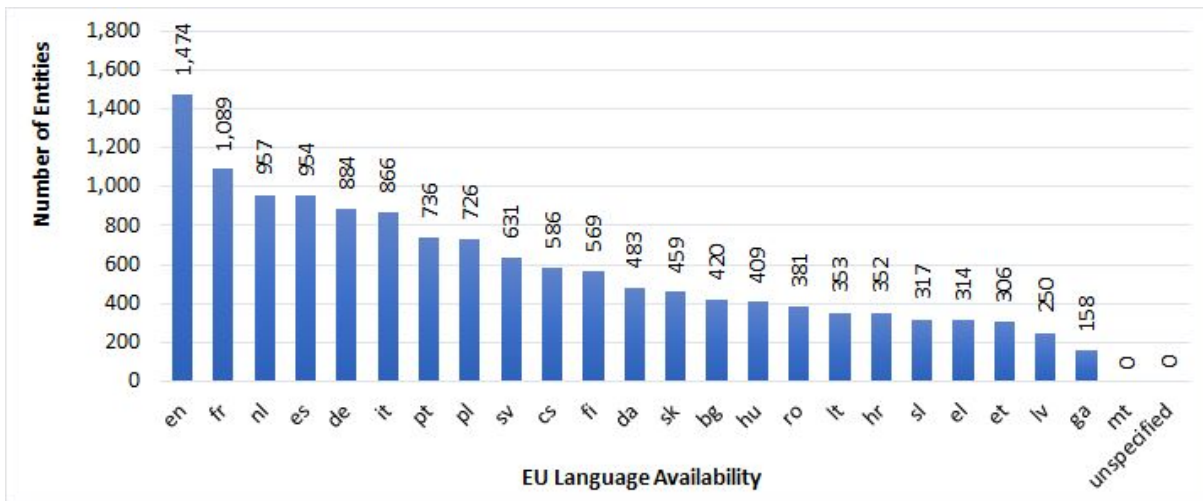
Concepts: 1,572 entities



Multilingual availability (skos:prefLabel) for Concepts (1,572 entities)

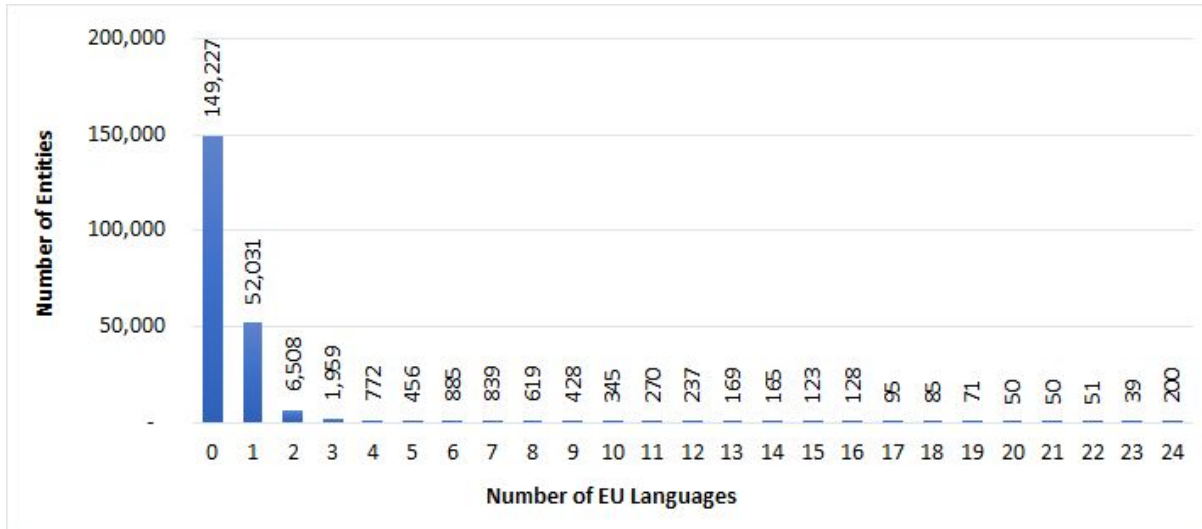
Each “Concept” entity have labels in 15.30 languages in average (SD=7.11, min=1, max=24, median=17). As shown in Figure 11, 13 entities do not have any labels in the EU languages. These entities instead have labels in other languages, such as Armenian, Arabic, Chinese, Japanese, Russian, Georgian, and Belarusian.

As shown in Figure 12, 93.77% entities have English labels. More than half the entities are also available in French, Dutch, Spanish, German, and Italian.



Multilingual availability (skos:prefLabel) for Concepts (1,572 entities) per language


Places: 215,802 entities



Multilingual availability (*skos:prefLabel*) for Places (215,802 entities)

In the “Places” EC, a large number of collections contain the “*skos:prefLabel*” property but no languages were specified for the values (shown in Figure 13). Therefore, these labels are not counted as the EU languages, resulting in a high proportion of entities with no labels in the EU languages (69%). On average, each entity has labels in 0.59 EU languages (SD=1.8, min=0, max=24, median=0).

The high proportion of unspecified-language “*skos:prefLabel*” only exists in the “Place” entities. One possible reason is that place names are often spelled the same way in most EU languages. An initial analysis further shows that language-specific “*skos:prefLabel*” were only used to list place names that were spelled differently in the particular language. These findings suggest that the multilingual availability of “Place” entities is higher than the Figure 13 and Figure 14 show, because the labels in the “unspecified” language may represent the correct values for most EU languages.



Multilingual availability (*skos:prefLabel*) for Places (215,802 entities) per language

A.3 Proportion of CH entities (from other Linked Data sources) that are available in the EC

To calculate the EC coverage, we compared the entities in the EC to Wikipedia and other linked data resources (Wikidata and ULAN).

We evaluated the EC coverage against a Wikipedia page containing a list of 3,213 painters²¹. For each person, we searched the Entity Collection using exact string-matching. 1,728 (54%) were found in the Entity Collection, while the remaining 1,485 (46%) were not available²². An initial analysis was performed on the first 10 unavailable entities. Five of these artists were American, four were European, and one Asian. They were either painters, or restorators. Four of these artists have relevant items in Europeana, however, were not included in the EC. More analysis is needed to include a bigger sample of the unavailable artists to gain better understanding of the EC coverage.

The coverage of the EC was also evaluated against other linked data resources. In this evaluation, we created test sets containing names of persons and concepts related to the cultural heritage domain²³. Wikidata and ULAN were chosen in this experiment because they contained list of relevant Agents in the CH domain. To build the test sets, we first selected a number of profession types that are relevant to the CH domain, such as “Artists”, “Visual Artists”, “Painters”, “Designers” and “Fashion Designers”. SPARQL queries were then created to extract the relevant entities of these chosen profession types in the Wikidata site²⁴ and the Getty site²⁵ for ULAN entities. The list of queries are shown in Appendix 1. List of relevant concepts in AAT are also extracted using the Getty site. These entities were then used as a test set to evaluate the EC coverage against.

In the future, the EC will also include “Works”, e.g., paintings. Therefore, a test dataset has also been created to include list of paintings extracted from Wikidata. This dataset can be used in the future to evaluate the coverage of the “Works” EC once this is available.

We report the coverage results of EC against these test sets in Table 5. The results show that EC coverage for Agents ranges between 8% to 31% compared to ULAN and Wikidata. A comparison to AAT shows that only 1% of AAT Concepts were available in the EC. These findings suggest that the contents of the EC can be further enriched using information from ULAN and Wikidata.

²¹ https://en.wikipedia.org/wiki/List_of_painters_by_name

²² The full list of names, including their corresponding entry in the EC (if found) is available in: <https://docs.google.com/spreadsheets/d/1W8lg8lCoktwm9bASMcxapcwWM5Wu22CA0tOrytK24nQ/>.

²³ The test sets are available for download in: <https://basecamp.com/1768384/projects/5774755/messages/76013887>

²⁴ <https://query.wikidata.org/>

²⁵ <http://vocab.getty.edu/queries>

Entity Types	Test Set	EC Coverage
Agents	ULAN: Artists	10% (15,811 of 162,513)
	Wikidata: Artists	23% (4,124 of 18,255)
	ULAN: Visual Artists	8% (16,769 of 217,814)
	ULAN: Painters	13% (9,629 of 72,810)
	Wikidata: Painters	20% (22,868 of 112,941)
	ULAN: Designers	9% (910 of 9,746)
	Wikidata: Designers	12% (266 of 2,255)
	ULAN: Fashion Designers	31% (58 of 186)
	Wikidata: Fashion Designers	26% (430 out of 1,676)
	Concepts	AAT: Concepts

Coverage evaluation

Appendix B: evaluation of linkages between Europeana objects and the Entity Collection

We report the initial results in evaluating the CHO-entities linkages:

Number of CHOs that contain links to Concepts²⁶	13 275 736 of 51 459 305 (25.8%)
Number of CHOs that contain links to Places²⁷	13 684 525 of 51 459 305 (26.59%)
Number of CHOs that contain links to Agents²⁸	1 400 248 of 51 459 305 (2.72%)

²⁶ Retrieved using the query:
https://www.europeana.eu/api/v2/search.json?wskey=api2demo&query=skos_concept:*http*data.europeana.eu* (19 April 2018)

²⁷ Retrieved using the query:
http://www.europeana.eu/api/v2/search.json?wskey=api2demo&query=edm_place:*http*data.europeana.eu* (19 April 2018)

²⁸ Retrieved using the query:
https://www.europeana.eu/api/v2/search.json?wskey=api2demo&query=edm_agent:*http*data.europeana.eu* (19 April 2018)

Number of enriched CHOs total²⁹	24 302 507 of 51 459 305 (47.23%)
---	-----------------------------------

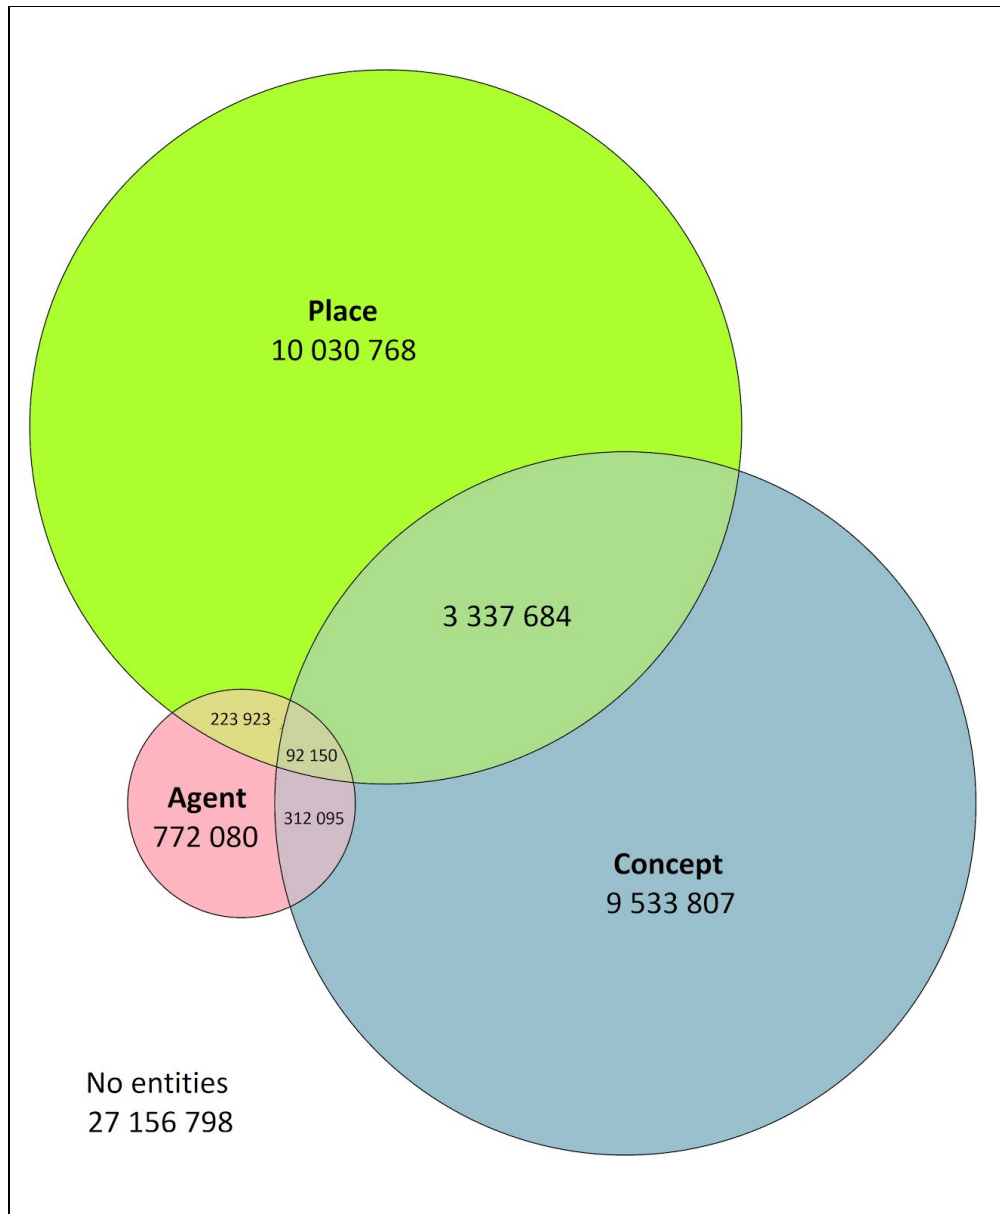
CHO-entity links evaluation

We also further analysed the proportion of CHOs that contain links to multiple entities, shown in Figure 15. The figure shows that a small number of CHOs (92 150; 0.18%) had links to all three entity types (Agents, Places, and Concepts) in the EC. The number of entity links to Agents is much lower compared to other entity types. A first reason is that the number of agents in the EC KG is much smaller - agents have been added much later in Europeana's automatic enrichment and the database of agents had been built carefully in order to avoid wrong enrichments. In domains such as Fashion and Photography, many items had information relating to the agent, e.g., fashion designers or photographers, but these entities did not exist in the EC. It is actually questionable, whether one can find suitable resources that describe such agents, as these domains have the characteristics of a "long tail" with many, lesser-known people involved in the creation of CHOs. These concerns should be further investigated in the context of the efforts for curating the EC KG³⁰. Another possible reason is that a large number of items in Europeana did not have any relevant Agents information. E.g., items such as photographs of the World War did not have any information about the creator (photographer).

As shown in Figure 15, a larger proportion of CHOs contain links to Places (26.59%) and Concepts (25.8%). The overlap between them, however, can be improved further. At the moment, only 6.49% (3 337 684) CHOs contain Places and Concepts links. Since the Concept entities should be relevant for all CHOs, future work should investigate ways to improving the availability of Concept entities for the remaining CHOs. This approach will also reduce the number of entities that did not have any entity links (currently at 27.16M entities; 52.77%).

²⁹ Retrieved using the query:
[https://www.europeana.eu/api/v2/search.json?wskey=api2demo&query=\(edm_place:*http*data.europeana.eu*%20OR%20skos_concept:*http*data.europeana.eu*%20OR%20edm_agent:*http*data.europeana.eu*\)](https://www.europeana.eu/api/v2/search.json?wskey=api2demo&query=(edm_place:*http*data.europeana.eu*%20OR%20skos_concept:*http*data.europeana.eu*%20OR%20edm_agent:*http*data.europeana.eu*))

³⁰ <https://docs.google.com/document/d/1A5Rb3Oe9edin5gdRppqFILLR0YPUodVOel3SdcBP00dA/>



Number of CHOs with links to different entity types

We also calculated the accuracy (precision) and the coverage (recall) of existing CHO-Entities linkages on a small set of evaluation documents. This initial dataset contain 10 documents randomly selected from the following Europeana collections: “Arts”, “Fashion”, “Music”, “Photography”, and “World War I” (see Table 7). All these documents are selected using the EN portal. This initial study helps to identify several issues that affect the qualities of these linkages and identify areas that require improvement. After this has been completed, a larger dataset should be used to gain a better insight of the qualities of the CHO-Entities linkages.

Firstly, for all the selected documents, we identified all the relevant “Agent”, “Concept”, and “Place” entities. These relevant entities were identified using the CHO metadata (if available) and by using further research if necessary. E.g., for a CHO item describing Vermeer’s “The Milkmaid” painting³¹, the following entities were annotated with the following entities:

- Agent: “Johannes Vermeer”³² (the painter)
- Concept: “Painting”³³, “Oil painting”³⁴
- Place: “Netherlands”³⁵

We then calculated the accuracy (precision) and the coverage (recall) of the current entity links against the annotated documents³⁶. E.g., if the current CHO item of “The Milkmaid” only contains 1 link to a concept called “Painting”, then P=1 (i.e., the only link is correct) and R=0.5 (i.e., because it only contains links to 1 out of the 2 relevant concepts identified in the gold standard). If no entities were linked by the CHOs, no numbers were computed. The results are reported in Table 7.

Item	Collections	Agent		Concept		Place	
		Prec	Recall	Prec	Recall	Prec	Recall
Item 1 ³⁷	Arts	1	0.5	1	0.33	-	0
Item 2 ³⁸	Arts	-	0	1	0.33	-	0
Item 3 ³⁹	Fashion	-	-	-	-	-	0
Item 4 ⁴⁰	Fashion	1	1	0.5	0.5	0	0
Item 5 ⁴¹	Music	1	1	-	0	0.14	0.5
Item 6 ⁴²	Music	-	-	-	0	-	0

³¹ https://www.europeana.eu/portal/en/record/90402/SK_A_2344.html

³² <http://data.europeana.eu/agent/base/60404>

³³ <http://data.europeana.eu/concept/base/47>

³⁴ <http://data.europeana.eu/concept/base/222>

³⁵ <http://data.europeana.eu/place/base/106>

³⁶ The annotated dataset (gold standard) is available in:

<https://docs.google.com/spreadsheets/d/16VNWelO0sHqi8pFYB4h00V6kpWW0rTMaeCd6ad1bGWg/edit?usp=sharing>

³⁷

http://www.europeana.eu/portal/en/record/2064116/Museu_ProvidedCHO_Nationalmuseum_Sweden_19494.html

³⁸ https://www.europeana.eu/portal/en/record/90402/SK_A_2344.html

³⁹ <https://www.europeana.eu/portal/en/record/2048224/00000137.html>

⁴⁰ https://www.europeana.eu/portal/en/record/2048232/europeana_fashion_61899095_jpg.html

⁴¹ http://www.europeana.eu/portal/en/record/2059209/data_sounds_11L0010327XX_0201.html

⁴² https://www.europeana.eu/portal/en/record/92085/lnb_zl_31628.html

Item 7 ⁴³	Photography	-	-	0	0	-	0
Item 8 ⁴⁴	Photography	1	1	-	0	0	0
Item 9 ⁴⁵	World War I	-	-	-	0	-	0
Item 10 ⁴⁶	World War I	-	-	-	0	-	0
Note: Precision refers to the accuracy of the existing links in the CHOs. Recall refers to the coverage of existing links evaluated against the gold standard data.							

Accuracy (precision) and coverage (recall) of links in the document test set

As shown in Table 7, the evaluation against the small evaluation dataset (gold standard) provides some insights on the current quality of the entity links. The results show that entities of “Agents” type contain high precision. In two items, however, the entities were not linked to the EC although both exist in the EC (one was linked to RijksMuseum URI instead).

In the gold standard, relevant “Concepts” entities were found that can describe most items, e.g., “black-and-white photography”, “world war 1”, “oil paintings”. However, most of the CHOs are not currently linked to any “Concepts” entities. The same situation applies to “Places”, where majority of the CHOs do not currently contain any links to “Places” entities. The findings shown in this results were reported using a very small dataset of 10 items. As a future work, more annotations should be gathered for a larger number of items to produce a better insights of the quality of the entity links.

Finally, the correct entities can only be annotated if they exist in the EC. I.e., limited coverage of the EC will directly affect the number of annotations in the gold standard. In this exercise, we found very limited terms in the EC to describe the items in the “Fashion” domain. None of the fashion concepts, such as “fashion show”, “dress”, “hat” were found in the EC. However, these concepts may be included in the EC in the future⁴⁷. Many fashion designers and photographers (from the “Photography” domain) were also not found in the “Agent” entities in the EC. The latter is to be expected due to the long tail of named entities in this area (i.e., those with lesser fame).

⁴³ https://www.europeana.eu/portal/en/record/92070/BibliographicResource_1000126223900.html

⁴⁴ https://www.europeana.eu/portal/en/record/2048429/item_YDNEDLESC5J5YTHBJUL254X2JE2WBO6J.html

⁴⁵ https://www.europeana.eu/portal/en/record/92085/lmb_zl_26249.html

⁴⁶ http://www.europeana.eu/portal/en/record/2051922/data_euscreenXL_3825790.html

⁴⁷ <https://docs.google.com/document/d/1A5Rb3Oe9edin5gdRppqFILLR0YPUodVOel3SdcBP00dAV/>