The Wind in Our Sails: Developing a Reusable and Maintainable Dutch Maritime History Knowledge Graph

Stijn Schouten
Victor de Boer
Vrije Universiteit Amsterdam
Department of Computer Science
Amsterdam, The Netherlands
stijnscho@gmail.com, v.de.boer@vu.nl

Lodewijk Petram
Huygens Instituut
Amsterdam, the Netherlands
lodewijk.petram@huygens.knaw.nl

Marieke van Erp
KNAW Humanities Cluster
DHLab
Amsterdam, the Netherlands
marieke.van.erp@dh.knaw.nl

ABSTRACT

Digital sources are more prevalent than ever but effectively using them can be challenging. One core challenge is that digitized sources are often distributed, thus forcing researchers to spend time collecting, interpreting, and aligning different sources. A knowledge graph can accelerate research by providing a single connected source of truth that humans and machines can query. During two design-test cycles, we convert four data sets from the historical maritime domain into a knowledge graph. The focus during these cycles is on creating a sustainable and usable approach that can be adopted in other linked data conversion efforts. Furthermore, our knowledge graph is available for maritime historians and other interested users to investigate the daily business of the Dutch East India Company through a unified portal.

CCS CONCEPTS

• **Applied computing** → **Arts and humanities**; • **Information systems** → **Graph-based database models**.

KEYWORDS

Knowledge Graph, Digital Humanities, Maritime history

ACM Reference Format:

Stijn Schouten, Victor de Boer, Lodewijk Petram, and Marieke van Erp. 2021. The Wind in Our Sails: Developing a Reusable and Maintainable Dutch Maritime History Knowledge Graph. In *Proceedings of the 11th Knowledge Capture Conference (K-CAP '21), December 2–3, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3460210.3493548>

1 INTRODUCTION

Recent large-scale digitization efforts have increased the volume of digitally available sources enormously [7]. This is a huge boon to digital humanities research, a multidisciplinary field, where researchers utilize computational processing power to develop and present insights that were not possible before in traditional humanities research [3]. However, using these digitized sources in large-scale research is not always straightforward. When data sets are available as standalone files, researchers have to go through

the laborious process of collecting the different files, interpreting and refining the data, defining relations between different entities, and finally creating the queries and filtering operations needed to answer their research questions. This can be a complex process; for one thing, interpretation of variables may differ between researchers [12] and the process is often repeated by every new researcher who wants to work with those data sets.

In this project, we create a knowledge graph of data related to the Dutch East India Company (Vereenigde Oostindische Compagnie, abbreviated as VOC). The VOC is often considered the first multinational company. It was a large Dutch trading company that was founded in 1602 and remained in business until the end of the eighteenth century [11]. The VOC kept records of its ships, their cargo, and the sailors and soldiers who ventured to Asia. This information is contained in four large data sets, three of which are available online from the websites of the Dutch National Archive (VOC Opvarenden, or VOC Crew members) and the Huygens Institute (Dutch Asiatic Shipping, Bookkeeper General Batavia and Places). The last data set was provided to us by the researchers from the Huygens Institute. The datasets are described in more detail in Section 4. Making these data sets available in a well-structured knowledge graph allows digital humanities scholars to accelerate research with these rich and unique data [15]. While many humanities researchers work with structured data such as databases and spreadsheets, linked data is still uncharted territory for many. This is partly due to unfamiliarity with the technology, and its at times steep learning curve to start using it [4].

The goal of this project is to design and develop a usable and sustainable knowledge graph, that will be hosted by the Huygens Institute, the research institute that manages three of the four data sets used in this project, and that aims to drive data-driven research. We adopt an iterative design process; in each iteration we capture additional knowledge from domain experts and the data sets. We are not the first to create linked data for historical resources (cf. [4, 16]) or even the first to do this for the historical maritime domain (cf. [8]) However, our work is the first to focus on sustainability of the knowledge graph, as well as in bringing together four Dutch historical maritime data sets.

The challenges in our design process are that the approach should be sustainable and that the designed artifact should be useful for researchers and domain experts. Measuring and determining the most sustainable approach is not straightforward and can be subjective. We therefore define four quality attributes that can be used to measure, steer, and lead the design to a sustainable solution.

This work is licensed under a Creative Commons Attribution International 4.0 License.

K-CAP '21, December 2–3, 2021, Virtual Event, USA.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8457-5/21/12.

<https://doi.org/10.1145/3460210.3493548>

These design requirements are set in cooperation with the project stakeholders and are as follows:

- *Transparency*: Clear and well-described design decisions and rules.
- *Accessibility*: Results and intermediate results should be accessible.
- *Re-usability*: Re-usable methods and semantic models.
- *Maintability*: Future-proof knowledge graph and infrastructure.

We evaluate our knowledge graph using competency questions following [1].

2 RELATED WORK

This project is on the intersection of computer science and humanities, therefore we discuss relevant related work from both disciplines, starting with best practices in knowledge modelling from the computer science/semantic web domain.

The design guidelines of developing an accessible, transparent, reusable, and maintainable knowledge graph are closely related to the FAIR [21] principles. Waagmeester et al. [19] provide us with tips and pointers on how to develop a knowledge graph according to the FAIR principles. Dimou [9] dives more into the technical side of knowledge graph engineering and inspires further research and comparison of these approaches. Finally, Jovanovik [17] brings these together in his linked data application development methodology, which focuses both on FAIR principles and the technical implementation.

Our work is not the first to design a knowledge graph for the historical maritime domain. De Boer et al. [8] started in 2014 with Dutch Ships and Sailors (DSS). The researchers involved in this project connected four different data sets in a first effort to unify the scattered Dutch historical maritime data. Additional work in this domain was carried out by Entjes [10] in 2015. Both projects score high in terms of *transparency* and *accessibility*, with design and conversion decisions being documented and reusable mapping scripts made available. In terms of re-usability, the verdict is less clear. DSS uses elements from multiple general-purpose ontologies and data models, and combines these with concepts and classes from a specific DSS model to describe the maritime activities. This design decision has definite upsides as it is generally recommended to apply widely used models and ontologies to enable interoperability with other graphs on the web.¹ The re-use potential within the (maritime) history domain is low, however: it is unlikely that the *dss:Ship* concept is used between projects outside of the DSS sphere.

As a use case for the OntoME platform Beretta et al. [2] developed another ontology for Dutch maritime data on ships and sailors in 2019.² OntoME is aimed at maintaining and distributing ontologies for digital heritage projects, and the approach of Beretta et al. therefore scores high in terms of maintainability. Furthermore, their ontology is based on the CIDOC CRM [6], a widely accepted and highly detailed model for digital cultural heritage. Concepts in this model, such as voyages, have a better chance of being re-used in other semantic web projects, thus offering great alignment possibilities. It is accessible and available for re-use on the OntoME

platform.³ The transparency and accessibility design principles of Beretta et al. and, more generally, the OntoME platform, leave room for improvement. Design decisions are not openly published, making it difficult for outsiders to understand and follow the thought process during the design.

Each of the discussed projects has its strengths and weaknesses. We adapt some of their strong aspects in our project, such as the transparent methodology of DSS and the model maintainability of OntoME.

3 RESEARCH METHOD AND DESIGN

3.1 Design decisions

In this project, the CIDOC CRM is selected as the ontology of choice. In contrast to the aforementioned DSS project, this will allow for further interoperability with other Digital Humanities Linked Data solutions. CIDOC CRM is a de facto domain standard and a safe choice; as it is for example already used in the OntoME project.

We use GraphDB for data conversion.⁴ GraphDB is transparent due to its interface; users can directly intuitively influence the mapping. It is reusable due to the ability to import and export mapping tables between projects. Maintainability also scores high due to the active development of the application. Other conversion tools were tested and examined as well. These include but are not limited to the tools of the Datalagend ecosphere⁵ which we found to be well capable of quickly producing high volumes of RDF data with minimal user input, but too rigid for modelling and linking entities. The RML-based editors, we investigated such as RMLEditor⁶ places a limit on file size and requests for an extension were not answered. Other RML editors (such as CARML,⁷ and Yarrml (matey)⁸) lack a workable visual or GUI representation. This downside makes the recommendation to use such an approach difficult for users outside computer and information sciences. Therefore, we deemed GraphDB the most suitable for our research goals.

3.2 Iterative design

We select *design science* as our central design methodology, as the primary goal of our project is the development of an artifact: the knowledge graph. In Design Science research, several research methods and frameworks exist (cf. [14, 20]). The main takeaway of these theories is the design-test cycle. In this iterative cycle, the researcher uses the designed artifact to solve the given problem while grounding the artifact in a socio-technical foundation. The artifact is evaluated and the evaluation results are used in the second iteration of the design-test cycle. The cycle continues until a certain predefined threshold is met.

We complete two design cycles. This allows for a review and evaluation period while still being time-efficient. In the first design-test cycle, we implement the approach of Jovanovik [17] with the goal of re-using some of the steps during the second design-test cycle. During the first design-test cycle we also conduct an open

¹<https://www.w3.org/standards/semanticweb/ontology>

²<https://ontome.net>

³<http://ontome.net/project/28>

⁴<https://graphdb.ontotext.com/documentation/standard/release-notes.html>

⁵<https://iisg.amsterdam/en/clariah>

⁶<https://rml.io/tools/rmleditor/>

⁷<https://github.com/carmml/carmml>

⁸<https://rml.io/yarrml/matey/>

and unstructured interview with a domain expert from the Huygens Institute. The domain expert is responsible for some and involved in most of the curation efforts of data sets used during in this project. An interview does not only help with understanding the data sets but also creates an understanding of the use case, how domain experts plan to use the artifact, and what they expect from it. Not all questions can be asked and answered during a single interview; some questions may arise during the cycle. Due to the pandemic, in person collaborations at the research institute's office were not possible; we therefore keep close contact via email. Finally, the documentation provided by the authors of the previous historical maritime conversions ([2, 8, 10]) is examined during desk research.

The first review period has two goals. First, errors, shortcomings, and other issues are investigated. Second, the review period enables the start and direction of the second design iteration. An unstructured discussion session with digital humanities experts and the domain expert completes the review period.

Finally, in the second review period, we test the design again and conduct a workshop with several stakeholders from the Huygens Institute and the University of Amsterdam. In this workshop, the developed artifact can be user-tested, thus validating the research results. Aside from the workshop, a short data story is written to illustrate the usefulness of the created knowledge graph (presented in Section 6).

We evaluate the *sustainability* aspect of our project through the design guidelines. *Usefulness* can be measured with the completion of the competency questions [1]. Our competency questions come from two sources: the domain expert of the Huygens Institute and a related project[10]. In the list of competency questions below, (D) indicates that the domain expert submitted the question; questions derived from the related project are denoted with (R):

- (D) Which VOC Chamber was accountable for the largest number of slaves transported?
- (D) How were the shipping routes in Asia divided between the various VOC chambers (for example, did ships from a specific chamber predominantly sail on certain routes)?
- (D) What was the average value of cargo on VOC return voyages per crew member/ship's ton, and how did this evolve over time?
- (D) To what extent was the value per ships' ton on return voyages correlated with the skipper's track record (how much experience, i.e., how many previous voyages / how quickly did a ship get to Asia on an outbound voyage)?
- (R) How did the price of goods change over time?
- (R) Perform a search for specific information in large data sets, such as names;
- (R) Can external factors, such as war, be linked to the intensity of shipping?

4 FIRST ITERATION

As discussed in the methodology section, we derive the structure of the iterations from the work of Jovanonik [17].

4.1 Data sets

We include three VOC related data sets in the first design iteration:

- *Dutch-Asiatic Shipping (DAS)* This data set contains information about all voyages from the Netherlands to Asia and back made by the VOC and its predecessors. Some core details are available for every voyage, such as the VOC Chamber that outfitted and contracted the ship, and the places and dates of departure and arrival.⁹ During these voyages, events occurred, such as a storm or a capture of an enemy vessel. These events are stored in the voyParticulars column. Voyages had a master and were made by ships that have a name, type and tonnage, and were built at shipyards. Ships could make multiple journeys and sometimes have their name changed. DAS contains the information of 4,700 voyages from the Netherlands to Asia and of 3,400 return voyages.
- *Bookkeeper-General Batavia (BGB)* In Batavia (now Jakarta in Indonesia), the bookkeeper and his clerks kept track of all imports and exports over sea. For a sizable portion of the eighteenth century, these records have survived. They were entered into a database in the past decade and published online.¹⁰ Similarly to the voyages in DAS, voyages in BGB leave and arrive on a specific date to and from a specific port. There are also overlapping voyages between the different data sets. These overlapping voyages were identified by previous research. The voyages transported several goods with a specific value, unit of measurement, and specification. Voyages partook on ships that have different names over time. Note that in comparison to DAS, less information about ships is captured in this data set. BGB consists of 55 volumes of the 18th century containing 18,722 records.
- *Places* The Places data set is a single table containing information about geographical places related to DAS and BGB. It contains toponym attestations from DAS and BGB that are linked to a standardized place name (toponym followed by a space and the ISO 3166-1 alpha 2 country code). GeoNames and Wikidata URLs, as well as geographical coordinates are available.¹¹ The places dataset consists 4,252 instances.

4.2 Data and ontology modeling

In the data and ontology modeling phase, we aim to map data concepts to our base ontology (CIDOC CRM) and add new attributes to that model if one of these concepts does not fit the model.

We initially identify four main components: voyages, ships, persons, and companies. Their relations include, but are not limited to, a voyage taking place on a ship, a sailor participating on a voyage, and the VOC chamber outfitting the ship. With these high-level concepts defined, we can specify the lower-level concepts.

Voyages form the starting point of the lower-level, detailed concepts. A voyage is the journey that a ship undertakes, moving from one place to another. In this journey, sailors, soldiers, and other personnel are employed to ensure that the ship's cargo safely reaches its destination.

We aim to re-use CIDOC CRM classes and properties where possible. If no good fit is available, we re-use classes and properties of Beretta et al. [2]. Finally, we create new classes and properties

⁹<http://resources.huygens.knaw.nl/das> Retrieved on: 13-05-2021

¹⁰<https://bgb.huygens.knaw.nl> Retrieved on: 13-05-2021

¹¹<https://www.geonames.org/>; <https://www.wikidata.org/>

Figure 1: Example of mapping branch

accommodate the envisioned data model. These newly created properties and classes are integrated into the CIDOC CRM model as a subclass of existing classes, thus making them interoperable with other CIDOC CRM extensions and models.

For reasons of readability, we do not display the entire ontology model, but instead show a branch starting from a voyage in Figure 1. The complete mapping figures can be found on the GitHub.¹²

We initially defined the **voyage** instance as a subclass of *CRM:E7 Activity*. We later changed this definition to a subclass of *CRM:E9 Move* to accommodate the linking of goods (physical objects) to the voyage via the property *CRM:P25 moved*. We define **ships** as a subclass of *CRM:E22 man-made object*. The property *CRM:P8 took place on or within* is used as property between the **voyages** and the **ships**. (Incidentally, the property description of *CRM:P8 took place on or within*, even uses this relation as an example: ‘It describes a period that can be located with respect to the space defined by an *CRM:E19 Physical Object* such as a ship or a building.’) We define the **name of the ship** through the property *CRM:P1 is identified by*, which links the ship instance to a *CRM:E41 Appellation* instance. Other definitions and classifications, such as ship types and productions, can be found in Figure 1.

4.3 Linked Data Transformation

In the transformation phase, we transform the acquired knowledge into RDF. In short, a table structure is mapped to the different positions within a RDF triple. We summarize our most important design decisions:

Time Time plays a central role in the maritime model. It relates to many concepts. Such as the arrival or departure of ships. In the initial phases of the conversion, the dates were sanitized and formatted according to the ISO 8601-1:2019

¹²<https://github.com/stijnschout3n/thewindinoursails>

standard, which does not allow for uncertainties. The reasoning behind this decision is as follows: more complex date formats are difficult to calculate with and can therefore lead to lower usability of the artifact. In consultation with the domain expert, we decided to add the original dates (which sometimes state uncertainties) as an extra attribute for researchers to consult while retaining data usability. Dates were formatted, according to CRM standards, as URIs and not as XSD:Date literals. This decision helps link events together that happened during the same time-span.

URIs We create URIs for unique (and unambiguous) entities. Identifying entities across different data sets is rather tricky[5] and it is not captured in this research project. Fortunately, other researchers have performed data resolution and linking on these data sets [13, 18]. For example, the BGB data set is linked at some entries to DAS voyages. In these cases, the DAS voyage entity is reused, effectively linking the two data sets. For goods shipped by the VOC, we create unique URIs for every good that was carried on a specific voyage. This is to allow for linking voyage-specific information (such as quantity or price) to the good.

4.4 Data interaction

After converting and connecting the different data sets, we make our knowledge graph accessible. We do this such that researchers are able to query the created knowledge graph and wrangle the retrieved data to achieve their research goal. Our knowledge graph can be accessed via <https://dutchshipsandsailors.nl/>.

During the project, we also experimented with the Metaphacts platform¹³ which is well integrated with GraphDB and comes in a pre-configured Docker container. Second, the Metaphacts platform

¹³<https://metaphacts.com/product>

hosts very usable graph exploration tools. These tools can help in answering the defined competency questions. Due to licensing issue, we cannot expose our data through Metaphacts, but invite readers to try out this platform, as well as exploring the data through the triplestore of their choice.

4.5 Review of the initial design

Modeling data sets into an existing ontology is a delicate process, even with high-level ontologies, such as CIDOC CRM. There is a continuous push and pull between modeling from a data perspective and modeling from an ontology perspective. In a perfect world, the ontology and the data are a match. In this case, there exists a class and property for every envisioned concept. Nevertheless, this is not the case when using a general ontology, as it defeats its purpose. Therefore, the envisioned concepts should be fitted within the scope of the ontology.

Differences in semantics between classes is not unique to this project. The domain experts foresaw this and started developing an extension for the CRM ontology. This extension, focusing on social, economic, and legal life, allows for modeling concepts outside of the museum artifact realm.¹⁴

We now take a step back to a more meta-level by checking the research questions' progress. Any forthcoming discrepancies can be fixed in the second design iteration.

Usability: The solution is approaching usable compared to the research goals. In Subsection 4.4, some of the competency questions could be answered with the interface. However, not all competency questions can be answered at this moment. The first design iteration did not include the very extensive VOC Opvarenden (VOCOP) data set. This data set contains rich records of the sailors and soldiers who crewed the VOC ships. These records will be included in the second design iteration, thus fulfilling the *usability* requirement.

Sustainability: This is not as strictly defined as the usability requirement. We can only review the approach and the decisions made during the first design iteration and determine where the sustainability can be improved. The approach, thus far, is *transparent* due to the extensive documentation. The triple store allows for an *accessible* environment where the knowledge graph can be accessed interactively. By reviewing different approaches and methods, the current approach seems reasonably *reusable* as well. As we use a CRM extension in the second iteration, we are actively promoting reusability.

5 SECOND ITERATION

In the second design iteration, we process the results of the first review period, and try to resolve any shortcomings regarding the research questions. This iteration will be noticeably compacter due to our improved understanding of the source data, making mere improvements rather than additions, and having fewer data sets to convert.

The design thus far consists of voyages (DAS), cargo moved on voyages (BGB), and the places where the voyages called (Places). However, several competency questions are related to the people involved in these activities. In the second design iteration, we add persons to the design to solve these competency questions.

The VOC kept extensive records of its contracted crew members ('Opvarenden' in Dutch). For the late seventeenth and eighteenth centuries, many pay ledgers have been preserved. Pay ledger entries include personal data, ranks, wages, and references to voyages. The voyages overlap with the voyages of the DAS data set. The information of 774,200 crew members was preserved in the VOC Opvarenden data set.

The core of the model does not change in the second iteration of the design. The changes in the second iteration are more apparent in the details. One of the more critical changes is related to the structure of the voyages. In summary of the previous iteration: a **voyage** is linked to a starting point via a departing activity and ends with the arrival of the voyage at its destination. Persons and goods are directly linked to the voyage; we assume that they were involved during the whole duration of the voyage. However, in the VOCOP data set, this is not the case. Persons could change ships during a stopover, effectively participating in two voyages. The first iteration could not adequately convey this concept.

To correctly convey this concept, we introduce a new class. The **leg** class is, similar to the voyage, a **phase**. It indicates a continuous part of the voyage from and to a port of call. A voyage can consist of multiple legs. The time between the end of the first leg and the beginning of the second leg is the stopover. Finally, the sailors are linked to the departing and arriving properties to indicate (dis)embarkment activities.

We choose the CIDOC CRM extension for economic, social, and legal life (SDHSS prefix) for mapping the VOCOP data set. To relate persons to a voyage, the *participation* concept is used. Other mappings include ranks and their descriptions.

The changes made during the second design iteration aim to improve the *reusability* of the artifact by utilizing an extended ontology. Creating mapping figures and tables improves *transparency* concerns. Converting the VOCOP data set ensures that the *usability* aspect as defined by the competency questions is met. To improve the *maintainability* aspect of the artifact, the OntoME platform is proposed. The power of OntoME is in the reuse of ontologies. The application has similarities with software development platforms such as Github in that it allows users to collaborate and discuss aspects concerning the creation and utilization of classes and properties. Furthermore, other researchers can easily reuse the captured knowledge within the ontology for their own projects.

Changes in the modeling have a cascading effect on the front-end part of the artifact. In the second design iteration, these issues are examined and fixed. Furthermore, the VOCOP data set was also integrated in the knowledge graph. An example of how the different sources can be explored in a single screen through the Metaphacts platform is shown in Figure 2.

Community interaction is not necessarily a research goal. However, it can have clear benefits. As [19] points out, sustainable growth of the knowledge graph can be realized with community involvement. A proposal that captures this issue is the utilization of the Geovistory platform.¹⁵ This platform can interact with a knowledge graph and hosts a text editor that can recognize entities and propose linkage to existing entities within the knowledge graph. Academics and citizen-scientists could for example, with approval,

¹⁴<http://ontome.net/project/64> Last visited: 16-09-2021

¹⁵<https://docs.geovistory.com/>

Figure 2: Visual integration of the four data sets

enrich the knowledge graph by browsing through the available data and adding connections within the graph.

6 EVALUATION

With the completion of the second design iteration, most of the issues discovered during the review period are resolved. These include, but are not limited to, better utilization of the CIDOC CRM model and its extensions, ontology management, community involvement, and the conversion of data on crew member (VOCOP). In the second review, we explore these improvements and related them to the primary research goal. The review consists of a workshop with five domain experts from the Huygens Institute and the University of Amsterdam. The expertise of the domain experts covered maritime history, digital data management and data science, all with an affinity for (digital) humanities research.

In the workshop, multiple discussions were brought up to challenge design decisions made during the research project. These topics include, but are not limited to, the utilization of CIDOC CRM, the conceptual mapping of a voyage, and the proposed infrastructure. The feedback was overall positive but also provided constructive criticism and included recommendations to put more focus on data provenance in future iterations.

Initially, one of the goals of this project was to create a usable knowledge graph in the field of Dutch historical maritime data. Discussion on what makes a knowledge graph useful stems mainly from the perspective of maritime historians, as they are the primary stakeholders. They provided the competency questions which are

used to test and improve the artifact. The answer to the solvability of these questions can be found in Table 1.

However, the developed artifact is not only intended to be used by maritime historians. Other interested parties can be drawn in by their research of an event or a long-lost forefather who had a remarkable VOC career. Determining if the knowledge graph is useful for these users is challenging as there is no clear picture of who these people are and what their research questions entail. Nevertheless, we found one subject and their research question is explored in a short data story.¹⁶ In short, the data story briefly dives into the life of a sailor, Hendrik Prins. Quantitative and qualitative data analytics are performed in order to portray the personal life and career of Hendrik Prins. The developed knowledge graph makes the process of combining data sources straightforward and quick.

7 DISCUSSION AND CONCLUSION

In this paper, we presented a reusable and maintainable knowledge graph for Dutch maritime history. The knowledge graph was created using an iterative design and development process that involved the conversion and linking of four different data sets concerning the Dutch East India Company's business dealings. The resulting knowledge graph consists of more than 21 Million RDF triples. The dataset is available through an online triple store¹⁷

¹⁶<https://github.com/stijnschout3n/thewindinoursails/blob/main/Datastory.pdf>

¹⁷<https://semanticweb.cs.vu.nl/eots/>

Question	Answer
Which VOC Chamber was accountable for the highest number of slaves transported?	Can be answered by querying voyages that transported slaves
How were the shipping routes in Asia divided between the various VOC chambers (for example, did ships from a specific chamber only sailed on certain routes)?	Can be answered. Already a preconfigured component in Metaphacts
What was the average value of cargo on VOC return voyages per crew member/ship's ton, and how did this evolve over time?	Can be answered. Goods, participants and prices are available
To what extent was the value per ships' ton on return voyages correlated with the skipper's track record (how much experience, i.e., how many previous voyages / how quickly did a ship get to Asia on an outbound voyage)?	Can be answered with data research, the components are available in the knowledge graph. Data paper provides an example.
How did the price of goods changed over time?	Can be answered, goods are linked to voyages which have dates
Search for specific information in large data sets, such as names.	Can be answered with search components of Metaphacts
Can external factors, such as war, be linked to the intensity of shipping?	If the external factors are captured in the same ontology, then yes.

Table 1: Competency questions including statement regarding solvability

and can be queried through a SPARQL endpoint¹⁸. The data itself is made available under open licences¹⁹. The data model and knowledge graph were evaluated using competency questions provided specifically for this project by a domain expert as well as competency questions formulated in earlier work for this domain. Furthermore, the model and knowledge graph were reviewed by five different experts with complementary knowledge regarding the domain and computational analysis in the context of digital humanities research.

The path for future work, appending more data sets and improving the model, is paved by the used methodology. By following the framework of Jovanovik [17], which focused on reusable components, other researchers can complete further iterations. Our documentation is made available to aid researchers in understanding the decisions made during the design. Other data sets could include the Northern Musterrols²⁰ or the Archangel and Elbing voyages²¹ in the Baltic sea and similar data sets from other countries.

It should also be noted that even with many changes in the second design iteration, the processing time of the second design iteration was much shorter than the first iteration. Whereas the first iteration took around six weeks to complete, we finished the second design iteration in two weeks. There is a steep learning curve in understanding the concepts, tooling, and methods used. Hopefully, other researchers will initiate new iterations and utilize the results of this research project to face a less steep learning curve.

There are limitations to our project. In the second design iteration, the proposed infrastructure is expanded with the use of OntoME and Geovistory. However, the two extensions are not integrated entirely as they are web-based and, therefore, difficult to integrate without direct support. In future iterations, the proposed architecture, including extensions, could be implemented in a production web server. However, this responsibility is outside the scope of this research project and lies with the curators and creators of the data.

During the workshop with domain experts, another limitation was discussed: provenance. Data-altering improvements can also be captured in the artifact. If these changes are stored in the knowledge graph, other researchers could backtrack and understand the changes made in the data, thus improving the transparency of the artifact. However, provenance information should be stored in a separate data layer to avoid graph pollution.

The developed artifact, a knowledge graph, aims to support historians and other users in their research. While the methodology also aims to develop the artifact sustainably, such that the artifact can be used and expanded over the years, we can conclude that the artifact can be deemed usable, as illustrated by Table 1. Close to all competency questions could be answered by using the knowledge graph. Due to the unavailability of data regarding events that had an impact on a voyage but did not on the ship, the last competency question ('Can external factors, such as war, be linked to the intensity of shipping?') could not be answered. Such research questions provide a useful insight into the types of research questions that domain experts are interested in and a solid direction for future extensions of the model and knowledge graph.

Furthermore, the participants of the workshop further confirmed the usability. A sidestep has been made by evaluating the artifact with the data story. During this evaluation, it was apparent that the

¹⁸<https://semanticweb.cs.vu.nl/eots/sparql>

¹⁹The Knowledge Graph and scripts to produce it is published as four separate datasets through the following DOIs: Places dataset: <https://doi.org/10.5281/zenodo.5509814>; VOCOPV dataset: <https://doi.org/10.5281/zenodo.5509820> DAS dataset: <https://doi.org/10.5281/zenodo.5507139>; BGB data <https://doi.org/10.5281/zenodo.5507126>

²⁰<https://www.noordelijkscheepvaartmuseum.nl/nl/collectie-en-onderzoek/onderzoek/monsterrollen>

²¹<http://resources.huygens.knaw.nl/archangel/app/inleiding>

artifact was useful for a broader range of people even though only a single evaluation was performed.

The results regarding the sustainability of the methods used and the developed artifact are multifaceted. There is no single answer that defines a sustainable artifact and process. In the related project section, we discussed the sustainability of the approaches taken by earlier project. We used these insights to shape our project in a more sustainable way.

Developing a sustainable knowledge graph consists of making design decisions. On a high or detailed level, design decisions influence the sustainability aspects. Picking between these design options can often lead to dilemmas where there is no clear answer. The choice between data usability and data correctness is one of these dilemmas. Do the data need to be accurate, even if this means the data will be difficult to process? In this project, we chose to store both original and cleaned dates. We faced another dilemma during the first review period. In this review, we examined the utilization of CIDOC CRM. The intention to re-use existing concepts of an ontology can conflict with the formal definition or use-case of that ontology. This in turn diminishes the sustainability as different projects might use the concept differently.

The two design cycles make the project stand out as the development of the artifact benefits immensely from receiving feedback. Design decisions made during the cycles are dependent on the researchers involved. Personal beliefs, experience, and other factors might steer the artifact in a different direction. Therefore, the transparent and accessible nature of the project is of utmost importance. It means that other researchers can not only continue with further iterations but also backtrack decisions and take a different path. Extended documentation, along with code and a Docker instance, can be found on GitHub.²²

ACKNOWLEDGMENTS

The authors want to thank Francesco Beretta, Vincent Alamercey for help with CIDOC CRM and participants in the workshop of KNAW HuC, the Huygens Institute and the University of Amsterdam.

REFERENCES

- [1] Kamal Azzaoui, Edgar Jacoby, Stefan Senger, Emiliano Cuadrado Rodríguez, Mabel Loza, Barbara Zdrasil, Marta Pinto, Antony J. Williams, Victor de la Torre, Jordi Mestres, Manuel Pastor, Olivier Taboureaux, Matthias Rarey, Christine Chichester, Steve Pettifer, Niklas Blomberg, Lee Harland, Bryn Williams-Jones, and Gerhard F. Ecker. 2013. Scientific competency questions as the basis for semantically enriched open pharmacological space development. *Drug Discovery Today* 18, 17–18 (Sept. 2013), 843–852. <https://doi.org/10.1016/j.drudis.2013.05.008>
- [2] Francesco Beretta, Vincent Alamercey, Sebastiaan Derks, Lodewijk Petram, and Jonas Schneider. 2019. Geohistorical FAIR data: data integration and interoperability using the OntoME platform. Time Machine Conference 2019. <https://halshs.archives-ouvertes.fr/halshs-02314003> Poster.
- [3] David M. Berry (Ed.). 2012. *Understanding digital humanities*. Palgrave Macmillan, Houndmills, Basingstoke, Hampshire ; New York. OCLC: ocn701020028.
- [4] Tobias Blanke, Gabriel Bodard, Michael Bryant, Stuart Dunn, Mark Hedges, Michael Jackson, and David Scott. 2012. Linked data for humanities research—The SPQR experiment. In *2012 6th IEEE international conference on digital ecosystems and technologies (DEST)*. IEEE, 1–6.
- [5] Peter Christen. 2012. *Data matching: concepts and techniques for record linkage, entity resolution, and duplicate detection*. Springer, Berlin ; New York. OCLC: ocn809643173.
- [6] Crm Cidoc. 2003. The CIDOC conceptual reference model.
- [7] <https://github.com/stijnschout3n/thewindinoursails>
- [8] Claire Warwick, Melissa M. Terras, and Julianne Nyhan (Eds.). 2012. *Digital humanities in practice: Digitisation and digital resources in the humanities*. Facet Publishing in association with UCL Centre for Digital Humanities, London. OCLC: ocn707962666.
- [9] Victor de Boer, Matthias van Rossum, and Rik Hoekstra. 2015. The Dutch Ships and Sailors Project. *DHCommons Journal* 1 (July 2015).
- [10] A. Dimou. 2020. *Knowledge Graphs and Big Data Processing*. Lecture Notes in Computer Science, Vol. 12072. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-030-53199-7>
- [11] Jeroen Entjes. 2015. *Linking Maritime Datasets to Dutch Ships and Sailors Cloud - Case studies on Archangelsvaart and Elbing*. Ph.D. Dissertation. VU Amsterdam. http://www.victordeboer.com/wp-content/uploads/2015/08/jeroen_entjes_final_thesis.pdf
- [12] Oscar Gelderblom, Abe de Jong, and Joost Jonker. 2013. The Formative Years of the Modern Corporation: The Dutch East India Company VOC, 1602–1623. *The Journal of Economic History* 73, 4 (Dec. 2013), 1050–1076. <https://doi.org/10.1017/S0022050713000879>
- [13] Bernhard Haslhofer, Antoine Isaac, and Rainer Simon. 2018. Knowledge Graphs in the Libraries and Digital Humanities Domain. *arXiv:1803.03198 [cs]* (2018), 1–8. https://doi.org/10.1007/978-3-319-63962-8_291-1 arXiv: 1803.03198.
- [14] B Hendriks, P Groth, and M van Erp. 2020. Recognising and linking entities in old dutchtext: a case study on voc notary records. *Collect & Connect Leiden* (2020).
- [15] Hevner, March, Park, and Ram. 2004. Design Science in Information Systems Research. *MIS Quarterly* 28, 1 (2004), 75. <https://doi.org/10.2307/25148625>
- [16] Lorna M. Hughes (Ed.). 2012. *Evaluating and measuring the value, use and impact of digital collections*. Facet Publ, London. OCLC: 774166752.
- [17] Eero Hyvönen et al. 2020. Linked open data infrastructure for digital humanities in finland. *Proceedings of Digital Humanities in Nordic Countries (DHN 2020)* (2020).
- [18] Milos Jovanovic. 2016. *Linked Data Application Development Methodology*. Ph.D. Dissertation.
- [19] Lodewijk Petram, Marijn Koolen, Rutger van Koert, Melvin Wevers, and Jelle van Lottum. [n.d.]. Data on the Maritime Workforce of the Dutch East India Company in the 18th Century. *to be published* ([n. d.]).
- [20] Andra Waagmeester, Gregory Stupp, Sebastian Burgstaller-Muehlbacher, Benjamin M Good, Malachi Griffith, Obi L Griffith, Kristina Hanspers, Henning Hermjakob, Toby S Hudson, Kevin Hybiske, Sarah M Keating, Magnus Manske, Michael Mayers, Daniel Mietchen, Elvira Mittraka, Alexander R Pico, Timothy Putman, Anders Riutta, Nuria Queralt-Rosinach, Lynn M Schriml, Thomas Shafee, Denise Slenter, Ralf Stephan, Katherine Thornton, Ginger Tsung, Roger Tu, Sabah Ul-Hasan, Egon Willighagen, Chunlei Wu, and Andrew I Su. 2020. Wikidata as a knowledge graph for the life sciences. *eLife* 9 (March 2020), e52614. <https://doi.org/10.7554/eLife.52614>
- [21] Roel J Wieringa. 2014. *Design science methodology for information systems and software engineering*. Springer.
- [22] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim Clark, Mercè Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo, Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J.G. Gray, Paul Groth, Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A.C. 't Hoen, Rob Hooft, Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Albert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos, Rene van Schaik, Susanna-Assunta Sansone, Erik Schultes, Thierry Sengstag, Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan van der Lei, Erik van Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg, Katherine Wolstencroft, Jun Zhao, and Barend Mons. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* 3, 1 (Dec. 2016), 160018. <https://doi.org/10.1038/sdata.2016.18>