CHAPTER 4
Approach

The advancement in the state of the art, illustrated in the previous part of this thesis, shows the complexity embedded in the act of interpreting art, and the attempts to entail such complexity with computational methods. The results of the survey on the availability of *icon* descriptions in current KGs show that few domain-specific data have been currently released, or, if present, the descriptive potential of ontologies is not fully exploited for formally expressing *icon* complexity. Nevertheless, to assess the efficacy of computational methods applied to the field, there is the necessity of thoroughly described, authoritative data on the topic.

The aim of the current research is to verify the feasibility and possible advantages of addressing iconological studies in a quantitative fashion, exploiting the expressiveness of the semantic web data structures, through the modeling of a LOD dataset of iconological interpretations. Therefore, this research addresses 1) a definition of an ontological modeling capable of representing the domain features of iconography and iconology, and 2) a verification of the advantages of addressing the domain of study with a quantitative approach. Consequently, the core contributions of the current research consist of representing iconological content in semantic data and assessing the usefulness of the availability of this content in a structured format.

4.1 Research questions, hypotheses and assumptions

As introduced in Section 1, the main research questions (MQ1) to which this research seeks to answer are the following:

- MQ1. How can an ontological modeling of *icon* interpretations represent the domain features to foster the access, analysis, and retrieval of *icon* content?
- MQ2. What would be the advantages of quantitatively analyzing a semantic network of *icon* interpretations?

Published in: Sofia Baroncini, Diving into artworks interpretations through the lenses of semantic data. An Application on Panofsky's Iconological Studies, Heidelberg: arthistoricum.net 2026, https://doi.org/10.11588/arthistoricum.1462

The first research question (MQ1) concerns the study of the domain (RQ1) and the ontological modeling (RQ2). The identification of the features of the domain defines how the topic is developed in the concrete steps to be taken to answer the research questions, namely: the ontology development, the dataset creation, and the analysis (see Figure 4.1).

The study of the domain is divided into two sub-questions:

- RQ1.1 What are the characteristics of the domain to be modeled?
- RQ1.2 What are the interests of the domain to be modeled? Which domain-specific research questions are addressed?

Following these two aspects of the domain, the ontological modeling is driven by two sub-questions:

- RQ2.1: What theoretical models are available in the domain-specific literature for describing the discipline? Among them, which one is feasible for ontological modeling?
- RQ2.2: What are the requirements that an ontology for the domain should include?

The latter main research question focuses on data analysis, which requires data creation (RQ3) and analysis (RQ4). In both cases, sub-questions related to the adopted theoretical approach and research questions are defined:

- RQ3. How to create data sources allowing us to quantitatively analyze *icon* interpretations?
  - RQ3.1. How can a dataset be structured according to a domain-specific theoretical model?
  - RQ3.2. What kind of data should be gathered to potentially reflect the interests of the domain?
  - RQ3.3. What further ontologies and vocabularies are needed for describing the aspects of the cultural objects and artworks involved?
  - RQ3.4 How this domain-specific data can be collected and processed?
- RQ4 What questions can be performed over such a dataset?
  - RQ4.1 What new questions can be addressed in such data modeled according to Panofsky's approach?
    - RQ4.1.1 Is it possible to quantitatively characterize an art historian's approach?
    - RQ4.1.2 Are there characteristics emerging from the data that can give us hints about the practically applied iconological method?

  - RQ4.2. Does data structured according to Panofsky's theory allow us to answer domain research questions?

Several hypotheses are formulated to address the research questions, which can be summarized as follows:

- Semantic web technologies are a feasible tool to express in a formalized way a complex domain of knowledge;
- Domain-specific theoretical approaches are suitable means for defining the characteristics of the domain, on which the ontological modeling can be based (RQ1);
- The interests of the domain can be extracted from art historians' scholarly articles and be expressed as research questions (RQ2);
- An art historian's claim is an interpretation, as multiple experts may have diverging, acceptable views concerning the same artwork (RQ2);
- The ontology efficacy in describing the domain can be successfully tested through 1) competency questions, 2) comparison with existing ontologies, and 3) automatic evaluation (RQ2);
- The possibility of performing domain-specific quantitative analysis can be verified with the aid of a domain-specific dataset (RQ3);
- The identified domain-specific research questions can be expressed as queries and performed over the dataset (RQ4).

Several assumptions underlie this research:

- Art historians express the results of their research in books and scholarly articles;
- Such texts are argumentative, and a degree of subjectivity is embedded in the art historian's claim;
- The data collected about the chosen case study are sufficiently free of error;
- The case study chosen constitutes a proof-of-concept of the possibility that the approach has in the domain of knowledge of iconography and iconology.

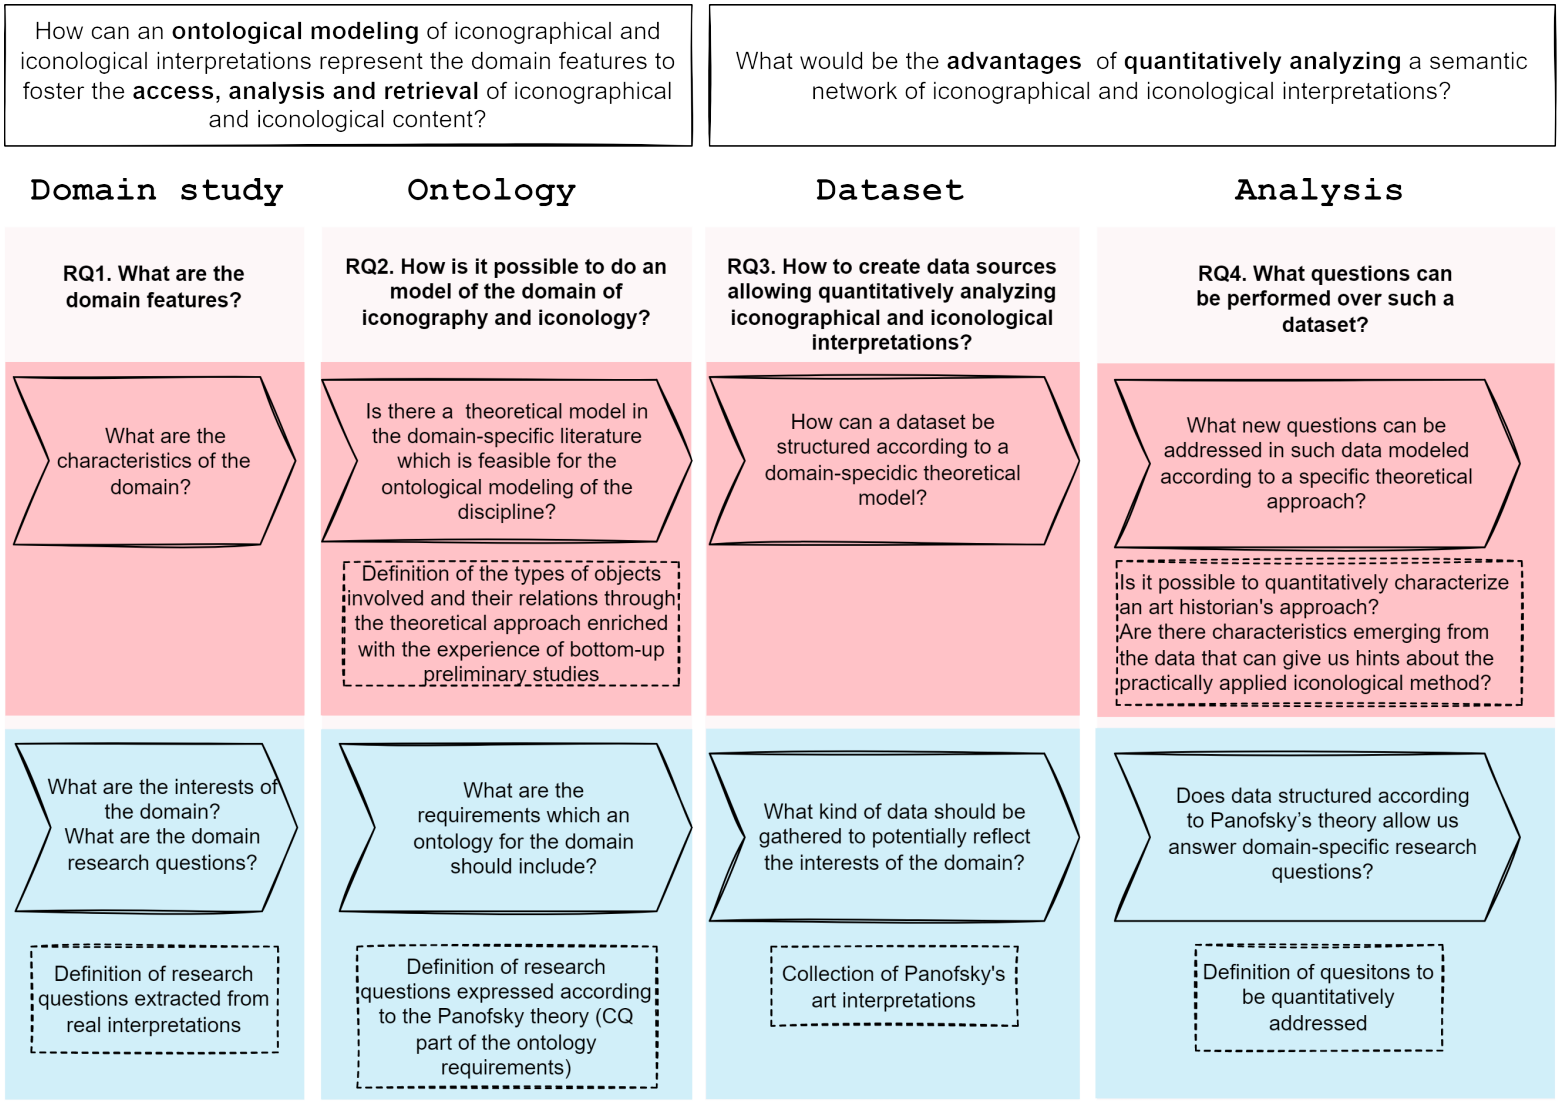

Figure 4.1: How the research questions are related

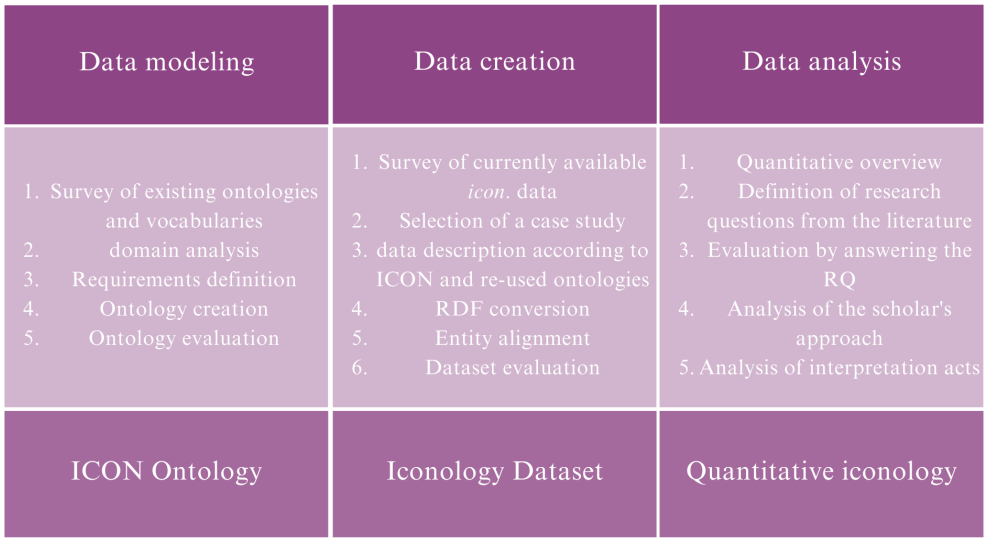

Figure 4.2: Approach overview

4.2 Approach to the research

To answer the research questions, we first studied the typological characteristics of the iconological field of study, then developed a feasible ontological model and manually collected and described data based on it.[1] Next, we conducted quantitative analyses. The approach is structured into three main parts: 1) domain study and ontological modeling, 2) data creation, and 3) data analysis, as shown in Figure 4.2.

4.2.1 Domain study

The domain study aimed to understand the characteristics and interests of an iconological approach by identifying the following: (1) the typological characteristics of the field through the top-down selection of the most feasible theoretical model available; and (2) the domain-specific research questions addressed by the discipline. These definitions will be fundamental to shaping the rest of the research, with the former being crucial for ontological modelling (Chapter 5) and the latter being essential for designing the analysis (Chapter 7).

1 The parts of the current Section relative to ontology and data anaylsis were respectively published in Sartini et al. (2023) and Baroncini, Daquino, and Tomasi (2023a) and reworked to be adapted to the current work

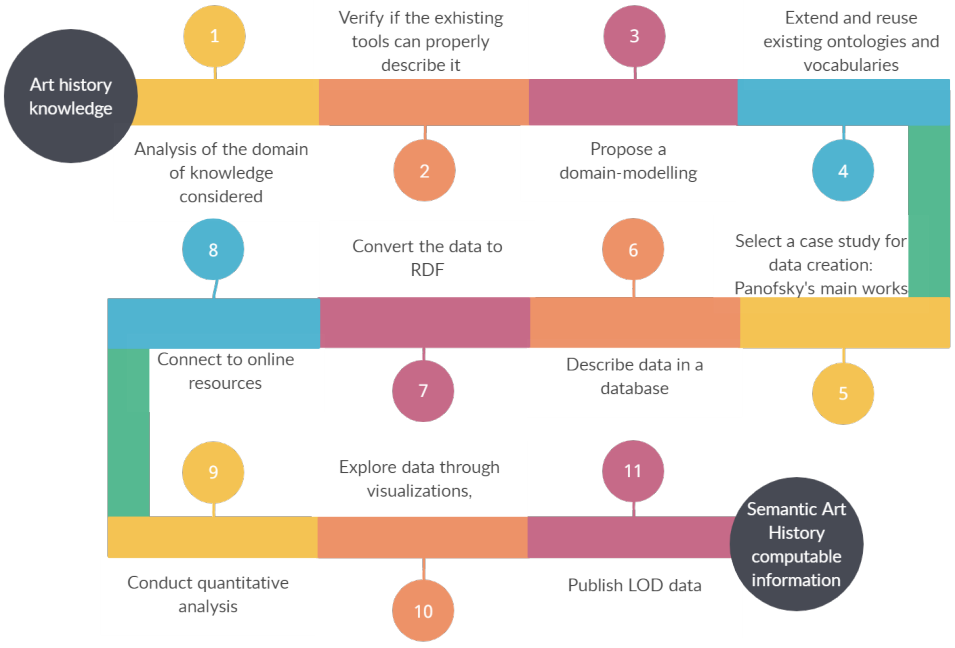

Figure 4.3: Timeline

4.2.1.1 Definition of domain characteristics

As illustrated in Chapter 2, several theoretical models were proposed by scholars to define the *icon* approach. Nevertheless, Panofsky's definition results in being a reference point for the domain and to be the most complete attempt of formalization. Therefore, we define the types of objects and relations involved in an *icon* study extracting them with a top-down approach from Panofsky's theory. To improve the definition of the domain characteristics, we enrich it with the outcomes of the preliminary study conducted over 11 case studies with a bottom-up approach (Baroncini et al., 2021). As a result of this section, we should obtain a brief description of the objects involved in an *icon* interpretation, the terminology, their role, and relations.

4.2.1.2 Definition of domain interests

With *domain interests*, we refer to the scope and aspects inquired by scholars during an *icon* interpretation. To this end, we employ a bottom-up approach, analyzing a selection of studies from diverse researchers to identify the typological questions they implicitly address in their analyses. This methodology is particularly emphasized by Szabo (2012), who conducted one of the first DAH studies focusing on the creation of databases designed for domain-specific, data-driven investigations into the study of artworks.

The research questions that were extracted have been organized thematically. These questions, which reflect the complex issues tackled by scholars, are considered in the subsequent steps of our approach. During the ontology development phase, they are framed according to Panofsky's theoretical perspective to define the Competency Questions. Following that, they are incorporated into the database creation stage. Ultimately, they are reformulated into subquestions to be quantitatively analyzed as part of both the data analysis and the dataset evaluation process.

4.2.2 Ontological modeling

Based on the analysis of the domain study, we formulate the requirements using the Ontology Requirements Specification Document (ORSD) and its guidelines presented in Suárez-Figueroa et al. (2009). These guidelines help to clearly define the scope, intended end users, terminology, and competency questions of the ontology.

The ICON ontology[2] was designed following the SAMOD (Peroni, 2016) and eXtreme Design (Presutti et al., 2009) methodologies. SAMOD is an agile methodology that focuses on the application of small iterative steps to model parts of an ontology. Each step is individually documented and combines motivating scenarios that derive from general domain descriptions with data-centric examples of descriptions formalized with the ontology. We re-use SAMOD methodology for the main part of the design, as we adopt the iteration-like structure and its outputs. The design process was divided into 4 SAMOD iterations, each dedicated to a particular aspect of the ontology. Each iteration contains a motivating scenario, a glossary with the definition of specific terms, a self-contained ontology prototype that contains only classes and properties relative to the corresponding iteration (with no references to external ontologies), the alignments to external ontologies, the aligned prototype, a series of competency questions formulated both in natural language and SPARQL (referring to the aligned prototype) and a Jupyter notebook that contains unit tests. All the competency questions were tested on real interpretations by Panofsky (1972) expressed using the ontology schema.

eXtreme Design is another agile methodology that divides the development of an ontology through iterations but focuses on the re-use of Ontology Design Patterns (ODP).[3] In fact, the methodology addresses modeling issues, either generic or related to the specific ontology that is being developed, with the re-use of modeling patterns. We specifically adopted this methodology when dealing with the re-use of ontology design patterns that were specialized in the context of our domain.

2 The ontology is available at https://w3id.org/icon/ontology/

3 ODPs are convceived as ontological fragments addressing common, recurring challenges encountered during the creation of ontologies. The patterns can be implemented in different ontologies so as to foster reusability and logic consistency (Gangemi & Presutti, 2009)

4.2.2.1 Ontology alignment

To promote the interoperability and reusability of the ontology, we connect to several external ontologies through alignment and reuse, following guidelines proposed by the state of the art (V. Carriero et al., 2020; Osman et al., 2021). To this aim, the ontology selection was guided by different principles: (i) standardization for CIDOC-CRM (Bekiari et al., 2021) and FRBRoo (Riva & Zumer, 2017) because they are considered standard frameworks in the domain, (ii) cognitive and formal analysis for the choice of DOLCE foundational ontology (Borgo et al., 2022; Gangemi et al., 2002) in its OWL version (DOLCE Zero), Simulation Ontology (Sartini et al., 2021), VIR (Carboni & de Luca, 2019), HiCO (Daquino & Tomasi, 2015) and CiTO (Shotton, 2010) as all offer design solutions to competency questions defined from the requirements in Section 5.2.

4.2.3 Dataset creation

To quantitatively address domain-specific research questions, a domain-specific knowledge graph should be created.

Since, to the best of the author's knowledge, structured data about *icon* interpretations are currently not available, and the complexity of an *icon* text of the scholarly literature does not allow automatic extraction of the content through automatic techniques (e.g., Entity Extraction), data had to be manually created.

For this study, we selected the interpretations by Erwin Panofsky as a case study and, therefore, as the target data to be gathered. For his relevance in the domain, and as his theory was selected for ontological modeling, his interpretations turned out to be a suitable candidate to perform detailed analysis on his own method. This choice allows us not only to perform quantitative analysis on authoritative domain-specific data but also to verify if his theory is consistently used in his own interpretations. By quantitatively characterizing Panofsky's approach, we assess the reliability of using his method in future applications. As his claims can not be considered a representative sample of the iconographical-iconological disciplines, which would require a selection of multiple art historians, the results of the analysis are limited to considerations about Panofsky himself and as a case study for the discipline.

4.2.3.1 Dataset creation workflow

The methodology adopted for the creation of datasets is based on the Best Practices for the Publishing of Open Linked Data published by the W3C.[4]

The dataset creation includes two main parts: 1) a database creation, and 2) a semantic dataset modeling and conversion. The first part concerns the design of a structure to manually collect the data to be subsequently converted to RDF format. Several reasons motivate the choice of adopting this pipepline in spite of creating a native RDF dataset. As the adopted modeling describes artworks with a high grade of granularity, multiple entities and reification classes are required for its description. To

4 https://www.w3.org/TR/2014/NOTE-ld-bp-20140109/#MODEL

the best of the author's knowledge, the currently available graph and semantic data storage platforms (e.g., Omeka-S[5], Neo4j[6]) are complex to customize for storing data with such a level of granularity, and the ease of data insertion can be affected by the complexity of the interface. Furthermore, since data are created manually, a tabular database gives the human compiler a better overview of the inserted data. Storing data in a custom-built database allows for 1) including multiple entities in the same table and 2) creating reification class URIs only during the conversion to avoid the time-consuming and error-prone task of creating such nodes by hand. Considering that our case is peculiar since the data had to be manually gathered, the database design is narrowly related to the ontological modeling decisions. For this reason, as shown in Figures 4.3 and 4.5, the core ontological modeling precedes the database design and population. Further database structure and ontological modeling improvement were introduced during the database population.

Consequently, the LOD dataset creation phases adopted from the W3C standard can be redefined as follows:

1. Ontological modeling (ICON ontology development and reuse of standard ontologies)
2. Data selection
3. Database design
4. Database population
5. Database cleansing
6. Definition of good URI creation rules
7. Definition of standard vocabularies reuse and alignment
8. Definition of a license
9. Data conversion
10. Provision of machine access to data

Whereas narrowly interlinked, for the sake of clarity, we describe in separate sections the phases concerning the database creation (i.e., database design, database population, database cleansing), semantic modeling and conversion (i.e., ontological modeling, definition of good URIs, standards reuse, data conversion), and technical requirements (license, provision of machine access).

5 https://omeka.org/
6 https://neo4j.com/

4.2.3.2 Database creation

Data selection As introduced, the study focuses on the art interpretations by Erwin Panofsky, chosen as a case study of the domain. Out of his bibliography, we selected three of his major and well-known books, plus one of his early articles. Although the selection of his own studies is not complete, it covers the main themes afforded by the art historian and can give a representative thematic insight into one of the core iconological themes of the Warburg school, namely the *Nachleben der Antike* in the art of the Renaissance and Middle ages.

Database design. We developed a tabular data structure that mirrors the ontology framework, incorporating information for each level of interpretation. The entities outlined correspond to the classes defined in the ICON ontology. Additionally, we included other relevant elements, such as metadata about artworks, individuals, books, and locations, in our database design. The artworks referenced in the preliminary study (Baroncini et al., 2021) served as a basis for testing and refining the database design.

Database population. We interpreted Panofsky's claims and described them according to the data structure. The quality of the data was assessed through value validation, which involved the utilization of controlled lists and spell checks. Since iconological interpretations do not have a fixed structure, no automatic recognitions were implemented to extract the knowledge. The entire process is based on the author's qualitative reading and interpretation of the text. Therefore, it should be considered that the statements obtained depend on the subjective comprehension of Panofsky's work by the author. For this reason, the author is indicated as responsible for the graph created, while Panofsky's text is always cited as the source of each statement.

Database cleansing. To avoid term inconsistencies, we used controlled lists to insert values and performed 1) spelling check and 2) detection of words used with both singular and plural forms. Before instantiating the conversion to RDF data, we verified that all the terms used were included in the controlled lists using a Python script.

4.2.3.3 Dataset creation

The **Ontological modeling** is driven by the results of the domain study. The definition of types led to the creation of the ICON ontology, with which we modeled the art interpretations collected. The code used for the creation, alignment, and evaluation of the RDF dataset is available on a GitHub repository.[7] The characteristics of the examples already used for the creation of the database were modeled to define the ontological modeling of other entities and features involved, not described by ICON.

7 https://github.com/SofiBar/IconologyDataset/

For this part, the approach of ontology reuse was adopted, since all characteristics could be described with available ontologies. The dataset modeling and conversion script were recursively refined during the database population phase to include relevant types of entities newly encountered.

Reused ontologies and vocabularies. During the dataset's modeling phase, we selected several ontologies to be reused to describe relevant aspects of the objects. CIDOC-CRM was selected as the reference standard for the Cultural Heritage community for the description of artworks and their features, books, and people. PRO[8] (Peroni et al., 2012) was reused for describing people's role in time, when relevant to the interpretation. Finally, CiTO relations were used for describing the information concerning the artistic interpretation (support, references, pieces of evidence).

Direct reuse or alignment was also performed at the vocabulary level. Types of artifacts, style, and period terms were described using Getty vocabulary AAT URIs. In all the other cases, an internal URI was created and aligned with one or multiple reconciled terms of controlled vocabularies. In detail, places and people names were aligned respectively to TGN or ULAN and VIAF. When possible, artworks were reconciled with Wikidata. The depicted subjects were then aligned with Iconclass, Wikidata, and HyperReal (Sartini et al., 2021). Types of relations between the artworks and the artifact (e.g. the relation between illumination and the manuscript page, or of a copy with the original) were modeled according to CIDOC-CRM relations, if available, and to Getty's CONA associative relations,[9] although their direct reuse was not possible, since they have not been released as Linked Data yet. The alignment will be described in more detail in Section 6.3.5.

The alignment was detected through a reconciliation script made for this purpose, based on a string similarity check,[10]. Results were manually checked and included in the main database (see Figure 4.4).

Good URI creation rules. We defined rules for creating good URIs during the conversion phase. When possible, we prefer to use *non-opaque URIs* by cleaning the items labels so as to remove special characters and white spaces. In other cases, a unique alphanumeric ID was assigned in the database and then converted to a URI. To better differentiate types of entities, we included the specification of their type in the URI. For fostering robustness and retrieval, the dataset itself was provided with a stable w3id prefix.

8 http://www.sparontologies.net/ontologies/pro

9 https://www.getty.edu/research/tools/vocabularies/guidelines/cona_3_5_associative_rels-copy.html

10 The code is available in the project's repository: https://github.com/SofiBar/IconologyDataset/tree/main/reconciliation

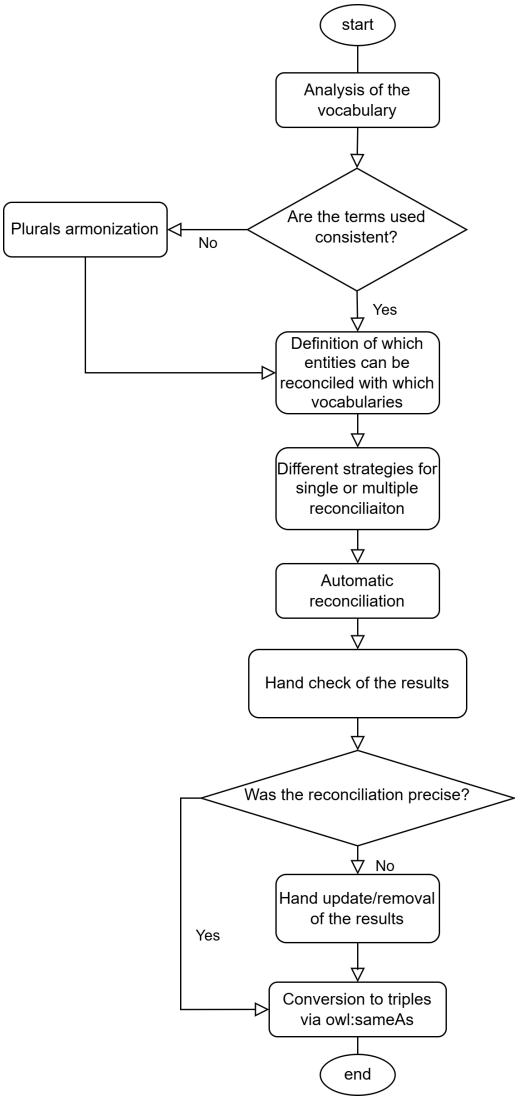

Figure 4.4: Workflow adopted for reconciliation

Data conversion. Data were converted to RDF format with a Python script using RDFlib,[11] a library for creating and querying RDF data. In the script, the database columns were mapped to the ontological modeling and converted into triples. The conversion was repeated several times, until 1) all the possible alignments were found and included in the dataset, and 2) all the errors were detected and corrected, as illustrated in the workflow in Figure 4.5.

4.2.3.4 Technical requirements

License. A Creative Common license (CC BY 4.0 SA) is included in the graph. Further metadata are provided, including author, source, language, and versioning. We reused Dublin Core Terms[12] and the PROV-O ontology[13] for this purpose.

Machine access to data. Data are exposed through a SPARQL endpoint[14] and provided with a unique identifier.[15]

4.2.4 Ontology and data evaluation

The dataset was evaluated according to a selection of metrics for the evaluation of data quality in Färber et al. (2018). Among them, the consistency metric, i.e., data compliance with the logical constraints established in the adopted ontologies, was evaluated through SHACL, which is, with ShEx, one of the languages aimed at describing and validating the consistency of RDF data against the logical constraints established by the used ontologies. This decision is supported by high relevance to the consistency validation measure provided by the literature for RDF data evaluation (Gayo et al., 2018), and specifically applied to Cultural Heritage data (Candela, 2023). SHACL was preferred to ShEx as it is the language recommended by the W3C for data consistency validation.[16] To verify if the dataset could address the domain-specific questions defined in the domain study, we defined subquestions that expressed the sought answer as a categorical value, subsequently expressed in SPARQL queries, which we performed over the dataset (Keet & Khan, 2024). We calculate the percentage of answered sub-questions and characterize answers as addressed, partially addressed (e.g. when limited data are available but the question could be answered via SPARQL queries), and not addressed. In doing so, our objective is to 1) assess the validity of the ontological model developed to represent Panofsky's theory and the dataset created for demonstration purposes, and 2) estimate to what extent quantitative analysis can be of help when answering traditional art history inquiries. The conformity of the dataset to the FAIR principles was assessed through self-questionnaires.

11 https://rdflib.readthedocs.io/en/stable/
12 http://purl.org/dc/terms/
13 https://www.w3.org/TR/2013/REC-prov-o-20130430/
14 https://projects.dharc.unibo.it/icondataset/sparql
15 https://w3id.org/icon/data/
16 https://www.w3.org/TR/shacl/

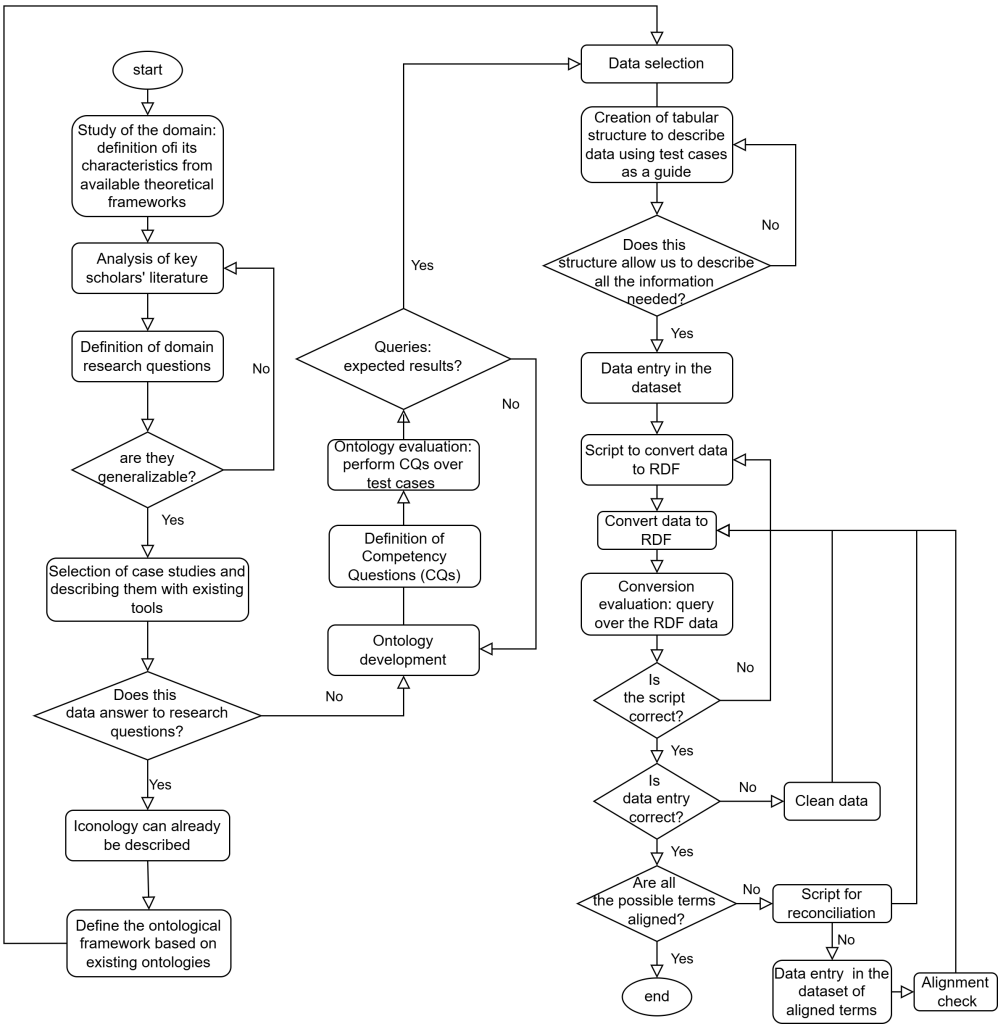

Figure 4.5: Workflow of ontological modeling and data creation

The ontology was validated across multiple dimensions. Its ability to represent and address key aspects of the iconography and iconology domains was evaluated by executing the competency questions outlined in the ontology requirements on a subset of the data set. Subsequently, a comparative analysis was performed against existing ontologies by modeling the same case study to evaluate the introduced innovations. Finally, an automated assessment of compliance with the FAIR principles was performed using the Foops tool (Garijo et al., 2021).

4.2.5 Analysis

This study uses quantitative analysis to gain insight into the benefits of using such an approach over a domain-specific dataset and to assess the further questions that can be addressed through data. The analysis, conducted over the above-described RDF dataset of iconological interpretations manually extracted from a selection of Panofsky's books, is structured in two main parts. We assess whether *icon* inquiries can be quantitatively addressed, by performing 1) the domain-research questions identified during the domain study (Section 5.1.2) as SPARQL queries over the database, and 2) what other questions could be addressed, with a focus on the art historical practice and methodology.

4.2.5.1 Iconographical and iconological research questions

The aim of RQ4.2 is to verify if and how many core research questions relevant to iconological studies can be answered in a quantitative way or, anyway, if a quantitative approach offers any valuable insights on the topic of interest. Such questions were obtained by rephrasing the questions extracted during the domain study phase into sub-questions, later expressed as SPARQL queries performed over the dataset. Whereas the rating of positive outcomes (that is, the percentage of questions successfully addressed) is presented in Chapter 8, Chapter 7 thoroughly illustrates the queries performed and their results.

4.2.5.2 New inquiries

We expressed hypotheses about what more questions could be addressed in the resulting data, which could not be answered with the traditional analysis, addressed in Section 7.3.

As the dataset is focused on Panofsky's studies, we suit the new inquiries drafted in Research Question 4.1. in his own practice. Therefore, more detailed questions can be expressed as follows:

- a. Can we characterise Panofsky's approach via data analysis?
  - RQ4.1a1. Does data analysis help us in confirming or refusing statements made about his theory?
  - RQ4.1a2. Does he consistently use the three levels of interpretation in his own studies?
  - RQ4.1a3. Can data modelled according to his theory fully represent the complexity characterising an iconological interpretation?
- b. Are there characteristics emerging from data that can give us hints on the practically applied iconological method? What are the elements needed for identifying a cultural phenomenon, in Panofsky's practice?

- b1. does the recognition of a deeper phenomenon need a thorough description at the previous levels?
- b2. are the artworks discussed in different books treated with the same level of detail?

Quantitative analysis of Panofsky's approach (RQ4.1a) As the dataset focuses on the case study of Erwin Panofsky, we perform analysis into Panofsky's method and research activity and validate the consistency of his method when applied to real-world attributions made by himself. To this end, we consider the following aspects: (1) whether he consistently reuses texts as primary sources of his interpretations, (2) whether he always uses a three-layer approach when motivating an iconological attribution, and (3) how sophisticated is the network of recognitions, cultural phenomena, and artworks considered in the analysis.

The role of texts (RQ4.1a1). Although it is a central reference for the formalization of the discipline itself, criticism was raised against Panofsky's theory, with several scholars claiming that he turned Warburg's iconological approach 'into a study of the transformation of artistic motifs through their interaction with texts'(Liepe, 2019, p. 17). Imdahl (2012) proposes to add an iconic level to the model, since, according to him, Panofsky's approach focuses on textual sources without giving the right relevance to the iconic language of artworks. Nevertheless, the centrality of documentary sources for conducting an *icon* analysis is underlined by Taylor (2008). In order to assess to which extent Panofsky motivates his interpretations with pieces of evidence and, in detail, textual sources, we analyze the number of interpretations based on a source, if text sources are more frequent in the recognition of a certain level, and of which type these are.

The three levels of interpretation coverage. RQ4.1a2 aims to represent the extent to which Panofsky adopts his own theory while interpreting artworks. Since the artworks in the dataset are described according to the three levels whenever the historian addressed them in the original text source, the approach to answer this question is to count how many artworks effectively have all the three levels described, and, if not, to retrieve the levels they are described with.

The complexity of an iconological interpretation. RQ4.1a3 examines how much the adopted interpretation theory is feasible to represent the complexity of an iconological interpretation. In fact, while explaining his own understanding, the art historian tends to cite artworks, subjects, and phenomena, with the risk of losing the interpretation unity when translated into data. One strategy which was adopted to maintain this unity is to use CiTO ontology and the network of citation relations (cito:citesAsEvidence, cito:givesSupportTo) to record the link between recognitions and artworks, especially when annotating third-level recognitions. In addition, cultural phenomena are often complex and an artwork can witness at the same time multiple phenomena, e.g. a contemporary phenomenon (e.g. the *interpretatio christiana* of classical deities as the

symbol of vices during the Middle Ages) and a more general one (the principle of disjunction, according to which every time a classical theme occurs, it is invested with moral meaning). Consequently, the coexistence of phenomena characterising artworks can give us a hint on the relations between phenomena themselves, and in the art historian's process of interpretation. Therefore, the analysis focuses on the network of recognitions, artworks, and phenomena, and on filtering of recognitions and artworks related by the CiTO relations. The network is created by including 1) all the phenomena in the selected network, 2) more artworks having a link to the same phenomena, and 3) other phenomena that these artworks are related to. In this way, the resulting network represents the maximum number of relations that can be retrieved from data. The analysis verifies whether the network presents a high number of connections, showing therefore the complexity of iconological interpretations.

Inquiries into iconological Panofsky's method RQ4.1b aimed at detecting which patterns emerge from Panofsky interpretations, which can give us hints on the applied iconological method. To this aim, we first verify whether the recognition of a deeper phenomenon needs a thorough description at the previous levels (RQ4.1b1), either in the artwork itself or in artworks presenting the same cultural phenomenon. Following, we address whether the style with which the art historian writes may affect the richness of the artworks description (RQ4.1b2) by relating the artworks to the books in which they are discussed.